

HIDDEN MARKOV MODELS FOR THE ACTIVITY PROFILE OF TERRORIST GROUPS*

BY VASANTHAN RAGHAVAN, ARAM GALSTYAN, AND ALEXANDER G. TARTAKOVSKY

University of Southern California

The main focus of this work is on developing models for the activity profile of a terrorist group, detecting sudden spurts and downfalls in this profile, and in general, tracking it over a period of time. Toward this goal, a d -state hidden Markov model (HMM) that captures the latent states underlying the dynamics of the group and thus its activity profile is developed. The simplest setting of $d = 2$ corresponds to the case where the dynamics are coarsely quantized as *Active* and *Inactive*, respectively. Two strategies for spurt detection and tracking are developed here: a model-independent strategy that uses the exponential weighted moving-average (EWMA) filter to track the strength of the group as measured by the number of attacks perpetrated by it, and a state estimation strategy that exploits the underlying HMM structure. The EWMA strategy is robust to modeling uncertainties and errors, and tracks persistent changes (changes that last for a sufficiently long duration) in the strength of the group. On the other hand, the state estimation strategy tracks even non-persistent changes that last only for a short duration at the cost of learning the underlying model. Case-studies with real terrorism data from open-source databases are provided to illustrate the performance of the two strategies.

1. Introduction. The increasing radicalization of divergent interest groups with different grievance profiles has led to an increased footprint of terrorist activity in many countries of the world. This threat has in turn led to an increased spending on counter-terrorism efforts that has shackled many of these economies (Mueller and Stewart, 2011; Belasco, 2006). In the light of these developments, it is critically important to track the activity of terrorist groups so that effective and appropriate counter-terrorism measures can be quickly undertaken to restore order and stability. In particular, detecting sudden spurts and downfalls in the activity profile of terrorist groups can help in understanding the dynamics of the terrorist group, thereby allowing the establishment to predict and even game the response to certain policy measures.

The basic element toward these goals is the collation of data on terrorist activities perpetrated by a group of interest. In this regard, it is well-understood that collecting terrorism

*This work has been supported by the Defense Threat Reduction Agency (DTRA) via grant #HDTRA-1-10-1-0086 and the U.S. Air Force Office of Scientific Research (AFOSR) via the MURI grant #FA9550-10-1-0569.

AMS 2000 subject classifications: Primary 62P25, 62C12; secondary 62L10

Keywords and phrases: hidden Markov model, self-exciting, terrorism, terrorist groups, Colombia, Peru, point process, changepoint detection

data relative to more common forms of crime is a challenging task (LaFree and Dugan, 2009). Despite these challenges, many databases such as ITERATE, the Global Terrorism Database (GTD) (LaFree and Dugan, 2007), the RAND Database on Worldwide Terrorism Incidents (RDWTI), etc., have made data on terrorism available in an open-source setting. However, even the best efforts on collecting data cannot overcome issues of temporal ambiguity, missing data, and attributional ambiguity that leads to mislabeled data. In addition, the very nature of terrorism makes it a rare occurrence from the viewpoint of model learning and inferencing. Thus, in addition to fitting observed terrorist data sufficiently accurately, an *ideal* model for the activity profile should: i) be motivated by a small set of hypotheses and remain robust to a certain looseness in the hypothesis claims, ii) be described by few parameters, and iii) be robust to missing or mislabeled data.

Prior work on modeling the activity profile of terrorist groups falls under one of the following three categories. Enders and Sandler (1993, 2000) use classical time-series analysis techniques to study both the short-run as well as the long-run spurt/downfall in world terrorist activity over the period from 1970 to 1996. On the other hand, LaFree, Morris and Dugan (2010) and Dugan, LaFree and Piquero (2005) adopt group-based trajectory analysis techniques (Cox proportional hazards model or zero-inflated Poisson model) to identify regional terrorism trends with similar developmental paths. The common theme that ties both these works is that the optimal number of underlying latent groups and the associated parameters that best fit the data remain variable. Further, the parameters are chosen to optimize a metric such as the Akaike Information Criterion (AIC), the Bayesian Information Criterion (BIC), etc., or via logistic regression methods. While a rich class of models ensures that reasonably acceptable model-fits can be arrived at, this is often at the cost of a large number of model parameters and with an *ad hoc* class that captures complicated dependency relationships between the endogenous and exogenous variables.

In addition, both sets of works take a *contagion* theoretic viewpoint (Midlarsky, 1978; Midlarsky, Crenshaw and Yoshida, 1980) that the current activity of the group is dependent on the past history of the group, which accounts for clustering effects in the activity profile of the group. This contagion viewpoint is put on a theoretical foundation in Porter and White (2012) where an easily decomposable two component self-exciting hurdle model (Hawkes, 1971; Cox and Isham, 1980) for the activity profile is introduced. Self-exciting models have become increasingly popular in as diverse fields as seismology (Ogata, 1988, 1998), gang behavior modeling (Mohler et al., 2011), insurgency activity that mirrors inter-gang violence (Lewis et al., 2011), epidemiology (Kim, 2011), etc. In its simplest form, the hurdle component of the model creates data sparsity by ensuring a pre-specified density of zero counts, while the self-exciting component induces clustering of data. However, the very nature of the self-exciting component implies that parameter estimation and inferencing algorithms are model-sensitive and have to be pursued on a case-by-case basis for different terrorist groups (and are hence not well-understood).

Contributions: In contrast to the self-exciting model, we hypothesize here that the cause of an increased/decreased frequency in attacks is some change in the group’s hidden states that reflect the dynamics of its evolution rather than the fact that the group has already carried

out attacks on either the previous day/week/month. In other words, while the self-exciting component of the model is best suited for earthquake activity (Ogata, 1988, 1998) where aftershocks are a direct consequence of current activity, or inter-gang violence (Mohler et al., 2011; Lewis et al., 2011) where attacks are retributive and often action-reaction consequences, activity modeling of terrorist groups is better motivated via a hidden Markov model (HMM) framework (Rabiner, 1989). We propose a d -state HMM for the activity profile where each of the d possible values corresponds to a certain distinct level in the underlying attributes. The simplest non-trivial setting of $d = 2$ with *Active* and *Inactive* states is shown to be a good model that captures most of the facets of real terrorism data.

We then develop a changepoint problem formulation to quickly identify a sudden spurt (or a sudden downfall) in the activity profile of the group. To address this problem, we propose two detection strategies. The first strategy is a stopping-time based on the Exponential Weighted Moving-Average (EWMA) filter that tracks the quantities of interest. This strategy does not require knowledge of the underlying model and is hence useful in settings where model learning is difficult or where model uncertainty is high. However, as a price for this luxury, only *persistent* changes (changes that last for a sufficiently long duration) can be detected with this strategy. The second strategy requires an estimate of the model parameters (that are learned over a training-period) to estimate the most probable state sequence using the Viterbi algorithm. We show that this strategy not only detects non-persistent changes, but also allows the identification of key geopolitical undercurrents (or events) that lead to sudden (major or minor) spurts/downfalls in a group's activity.

The advantages of the proposed framework are many. The proposed d -state HMM is built on a small set of easily motivated hypotheses and is parsimoniously described by $d(d + 1)$ model parameters. While observed data rarity can be explained by state transitions, clustering of attacks can be attributed to different intensity profiles in the different states. In addition to resulting in a good model-fit with real terrorism data, the proposed framework offers two significant advantages over existing approaches: i) the HMM paradigm allows the use of a mature and computationally efficient toolkit for model learning and inferencing (Rabiner, 1989), and ii) it offers a smooth trade-off between accurate modeling of terrorist activity and computational efficiency.

Organization: This paper is organized as follows. Section 2 illustrates some of the key qualitative features that describe the activity profile of terrorist groups and explains how prior work has modeled these features. In Section 3, a HMM framework is proposed to capture the key facets of the activity profile. Section 4 develops two strategies for detecting spurts and downfalls in the activity profile: a model-independent approach based on the EWMA algorithm and a HMM approach based on the Viterbi algorithm. The efficacy of the proposed approaches is studied with real datasets in Section 5. In Section 6, some of the key limitations in existing work that motivated the HMM framework are revisited and it is shown that the proposed approach overcomes these limitations. Concluding remarks are drawn in Section 7.

2. Preliminaries.




FIG 1. (a) Number of days of terrorist activity, and (b) Total number of attacks in a δ -day time-window as a function of time for FARC ($\delta = 15$ days). The key geopolitical events in this period are also marked.

2.1. *Qualitative Features of Activity Profile.* A typical example of the activity profile is presented in Fig. 1 where the number of days of terrorist activity in a $\delta = 15$ day time-window and the total number of attacks within the same time-window are plotted as a function of time. The focus of this example is *Fuerzas Armadas Revolucionarias de Colombia* (FARC), the subject of our focus in Section 5.1. The data for Fig. 1 is obtained from the RDWTI and corresponds to incidents over the nine-year period from 1998 to 2006. From Fig. 1 and a careful study of the activity profile of many terrorist groups from similar databases, we highlight some of the important features of terrorism data that impact modeling:

- **Temporal Ambiguity:** The exact instance (time) of occurrence of a terrorism incident is hard to pinpoint in reality. This is because accounts of most terrorist events are from third-party sources. Thus the granularity of incident reportage (that is, the time-scale on which incidents are reported) that is most relevant is *discrete*, typically days.
- **Attributional Ambiguity:** Further, in many of the databases, there exists an ambiguity in attributing a certain terrorism incident to a specific group when multiple groups contest on the same geographical turf. Some of this ambiguity can be resolved by attributing an incident to a specific group based on the attack signature (attack target type, operational details, strategies involved, etc.). However, this is an intensive manual exercise and there is necessarily a certain ambiguity in this resolution.
- **Data Sparsity:** Despite the recent surge in media attention on trans-national terrorist activities and insurgencies, terrorism incidents are “rare” even for some of the most active terrorist groups. For example, the data in Fig. 1 corresponds to 604 incidents over a nine-year period leading to an average of ≈ 1.29 incidents per week. While a case can be made that these databases report only a subset of the true activity, the

fact that significant amount of resources have to be invested by the terrorist group for every new incident acts as a natural dampener toward more attacks.

These three features make a strong case for simplistic models for the activity profile of terrorist groups that can be described with few parameters and are motivated by few hypotheses, in contrast to elaborate and expansive models of many parameters. Further, any model should be robust to small number of errors in terms of missing or mislabeled data.

2.2. Prior Work. The typical observable of a terrorist group is its activity profile with the observations being random point occurrences corresponding to terrorism incidents perpetrated by the group over a given time-period of interest. The previous discussion showed why a model for terrorism cannot be developed in continuous-time. A point process model in discrete-time is thus the most general framework (Cox and Isham, 1980) to model the activity profile.

Let the first and last day of the time-period of interest be denoted as Day 1 and Day \mathcal{N} , respectively. Let M_i denote the number of terrorism incidents on the i th day of observation, $i = 1, \dots, \mathcal{N}$. Note that M_i can take values from the set $\{0, 1, 2, \dots\}$ with $M_i = 0$ corresponding to no terrorist activity on the i th day of observation. On the other hand, there could be multiple terrorism incidents corresponding to independent attacks/targets on a given day reflecting the high level of coordination between various sub-groups of the group. Let \mathcal{H}_i denote the history of the group's activity till day i . That is, $\mathcal{H}_i = \{M_1, \dots, M_i\}$, $i = 1, 2, \dots, \mathcal{N}$ with $\mathcal{H}_0 = \emptyset$. The point process model is complete if $P(M_i = r | \mathcal{H}_{i-1})$ is specified as a function of \mathcal{H}_{i-1} for all $i = 1, \dots, \mathcal{N}$ and $r = 0, 1, 2, \dots$.

We noted in Section 1 that prior work on models for the activity profile fall under one of three categories. In the time-series techniques pioneered by Enders and Sandler (2000), a non-linear trend component, a seasonality (cyclic) component, and a stationary component are fitted to the time-series data of worldwide terrorism incidents. In particular, the following model-fit is proposed for $\{M_i\}$:

$$\sum_{i \in \Delta_q} M_i = \sum_{j=0}^n \alpha_j q^j + \beta \sin(\omega q + \phi) + \mu_q,$$

where Δ_q denotes the period corresponding to the q th quarter in the period of observation and $\{n, \alpha_0, \dots, \alpha_n, \beta, \omega, \phi, \text{Var}(\mu_q)\}$ are constants to be optimized over an appropriately chosen parameter space. This modeling effort results in an eight parameter model (4 for the trend component, 3 for the seasonality component, and a variance parameter for the stationary component) which is then used to identify a rough 4 and 1/2-year cycle in terrorism events. Further, a non-linear trend and seasonality component ensures that trends in terrorism cannot be predicted thus explaining the observed boom and bust cycles in terrorist activity.

In (Dugan, LaFree and Piquero, 2005), the hazard function $\text{Haz}(t)$ of the time to the

next hijacking (denoted as the random variable T) is modeled as

$$\begin{aligned} \log(\text{Haz}(t)) &= \log(\text{Haz}_0) + \beta_1 \cdot \text{Policies} + \beta_2 \cdot \text{Major Purpose} + \beta_3 \cdot \text{Context}, \\ \text{Haz}(t) &= \lim_{\Delta t \rightarrow 0} \frac{\text{P}(t \leq T < t + \Delta t)}{\Delta t}, \end{aligned}$$

where Haz_0 is a baseline hazard function, and $\{\beta_i\}$ are optimized from a certain model class. On the other hand, the hurdle model of [Porter and White \(2012\)](#) is described as

$$\text{P}(M_i = r | \mathcal{H}_{i-1}) = \begin{cases} e^{-(B_i + SE_i(\mathcal{H}_{i-1}))}, & r = 0 \\ \frac{r^{-s}}{\zeta(s)} \cdot (1 - e^{-(B_i + SE_i(\mathcal{H}_{i-1}))}), & r \geq 1 \end{cases}$$

where $s \in (1, \infty)$ is an appropriately chosen parameter of the Zipf¹ distribution, $\zeta(s) = \sum_{n=1}^{\infty} n^{-s}$ is the Riemann-zeta function, B_i is a baseline process, and $SE_i(\cdot)$ is the self-exciting component given as

$$SE_i(\mathcal{H}_{i-1}) = \sum_{j: j < i, M_j > 0} \alpha_j g(i - j)$$

for another appropriate choice of decay function $g(\cdot)$ and influence parameters $\{\alpha_j\}$. Note that the Zipf distribution falls in the class of power-law distributions (see Supplementary Part A) and is also explored in ([Clauset, Young and Gleditsch, 2007](#)). [Porter and White \(2012\)](#) study model-fits from a class described by eight parameters and show that a four parameter model optimizes an AIC metric for terrorism data from Indonesia/Timor-Leste over the period from 1994 to 2007. This model is shown to accurately capture terrorism data (especially the extreme outliers such as days with 36, 11, and 10 attacks).

While the above set of models reasonably accurately capture terrorism data, they are cumbersome with too many parameters to be estimated. Further, the choice of metrics such as AIC or BIC lead to the best parsimonious model-fit from a certain class of models for observed terrorism data. Nevertheless, it is more reasonable to assume that the cause of an increased frequency of attacks is due to some changes in the group's tactics (a hidden state), rather than the fact that the group has already carried out attacks previously. Thus, activity modeling of terrorist groups is better motivated via a HMM.

3. Proposed Model for the Activity Profile. Toward this direction of building a HMM for the activity profile, we make the following simplifying hypotheses in this work:

- **Hypothesis 1:** The activity profile of a terrorist group $\{M_i, i = 1, \dots, \mathcal{N}\}$ depends only on certain states (\mathbf{S}_i) in the sense that the current activity is conditionally independent of the past activity/history of the group given the current state:

$$\text{P}(M_i | \mathcal{H}_{i-1}, \mathbf{S}_i) = \text{P}(M_i | \mathbf{S}_i), i = 1, 2, \dots$$

¹This distribution is sometimes referred to as discrete Pareto.

In other words, these states completely capture the facets from the past history of the group in determining the current state of activity. While attaching specific meanings to \mathbf{S}_i is not the focus of this paper, an example in this direction is the postulate by Cragin and Daly (2004) that the *Intentions* and the *Capabilities* of a group could serve as these states.

- **Hypothesis 2:** It is clear that the dynamics of terrorism is well-understood if the underlying states $\{\mathbf{S}_i\}$ are known. However, in reality, \mathbf{S}_i cannot be observed directly and we can only make indirect inferences about it by observing $\{M_j, j = 1, \dots, i\}$. To allow inferencing of \mathbf{S}_i , we propose a simple d -state model that captures the dynamics of the group's latent states over time. That is, $\mathbf{S}_i \in \{0, 1, \dots, d-1\}$ with each distinct value corresponding to a different level in the underlying attributes of the group. Further, we penalize frequent state transitions in the terrorist group dynamics by constraining \mathbf{S}_i to be fixed over a block (time-window) of δ days, where δ is chosen appropriately based on the group dynamics.

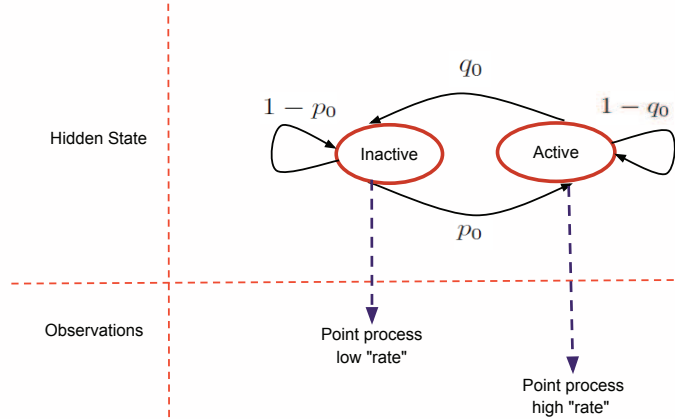


FIG 2. A simple $d = 2$ state HMM for the activity of a terrorist group.

A typical illustration of this framework with $d = 2$ is provided in Fig. 2 where the state over the n th time-window ($\Delta_n, n = 1, 2, \dots, K$) corresponding to $\Delta_n = \{(n-1)\delta + 1, \dots, n\delta\}$ and $K = \lceil \frac{N}{\delta} \rceil$ is given as

$$\mathbf{S}_i \Big|_{i \in \Delta_n} = s_n, \quad s_n \in \{0, 1\}.$$

In the *Inactive* state ($s_n = 0$), the underlying M_i form a low-“rate” point process, whereas in the *Active* state ($s_n = 1$), the M_i form a high-“rate” point process. Thus, a state transition

from $s_{n-1} = 0$ to $s_n = 1$ indicates a spurt in the activity of the group, whereas an opposite change indicates a downfall in the activity. This evolution in the states of the group is modeled by a state transition probability matrix $\mathbf{P} = \{\mathbf{P}_{ij}\}$, where

$$\mathbf{P} = \begin{bmatrix} 1 - p_0 & p_0 \\ q_0 & 1 - q_0 \end{bmatrix}$$

with

$$\mathbf{P}_{ij} \triangleq \mathbf{P}(s_n = j | s_{n-1} = i), \quad i, j \in \{0, 1\}.$$

In general, there exists a trade-off between accurate modeling of the group's latent attributes (larger d is better for this task) versus estimating more model (transition probability matrix) parameters (smaller d is better). While attention in the sequel will be restricted to the $d = 2$ setting because of the implementation ease of the proposed approaches in this setting, these approaches can be easily extended to a general d -state setting.

To model the observations in either state (that is, $\mathbf{P}(M_i | \mathbf{S}_i = j)$, $j = 0, 1$), a geometric density can be motivated as follows. Consider a setting where the group orchestrates M_i attacks on the i th day till the success of a certain short-term policy objective can be declared. Every additional attack contributes equally to the success of this objective and as long as the group's objective has not been met, attaining this objective in the future is not dependent on the past attacks. In other words, the group remains *memoryless* (or is *oblivious*) of its past activity and continues to attack with the same pattern as long as its objective remains unmet. A slight modification in the group dynamics that assumes that a group resistance/hurdle needs to be overcome before this *modus operandi* kicks in leads to a hurdle-based geometric model for $\{M_i\}$.

While these strategies may best describe a certain group, other groups could adopt different strategies. In fact, strategies could also change with time. Thus, a certain model could make more sense for a given terrorist group than other models. In our subsequent study, we consider the following one-parameter models with support on the non-negative integers: Poisson, shifted Zipf, and geometric. We also consider the following two-parameter models: Pólya², (non-self-exciting) hurdle-based Zipf, and hurdle-based geometric. Of these six models, the geometric and the hurdle-based geometric allow simple recursions for estimates of model parameter(s) via the Baum-Welch algorithm, while the shifted Zipf and the hurdle-based Zipf distributions allow heavy tails (see Supplementary Part A for details). Further details on these models such as the associated probability density function, log-likelihood function, Maximum Likelihood (ML) estimate of the model parameters, and a formula for the AIC are also provided in Supplementary Part A. The study of fits of these models to specific terrorist groups is undertaken in Section 6.2.

4. Detecting Spurts and Downfalls in Activity Profile. We are interested in solving the *non-causal* inference problem:

$$(4.1) \quad \widehat{s}_n \Big|_{n=1, \dots, K} = \arg \max_{s_1, \dots, s_K \in \{0, 1\}^K} \mathbf{P}(\mathbf{S}_i = s_n, i \in \Delta_n | \{M_j, j = 1, \dots, N\}).$$

²Also referred to as negative binomial with positive real parameter.

In particular, we are interested in identifying state transitions that correspond to either a spurt or a downfall in activity. For this, note that while the parameters associated with the model proposed in Section 3 can be learned from terrorism data, there is often an interest in approaches that are independent of these parameters, and hence not sensitive to the parameter estimation algorithms or the length of the training-period. Thus, we propose two approaches to address (4.1): the first strategy is independent of the model parameters, whereas the second strategy exploits the underlying HMM structure.

4.1. Model-Independent Strategy. We propose an algorithm for spurt detection based on the Exponential Weighted Moving-Average (EWMA) filter. The EWMA approach was first introduced by Roberts (1959) for (continuously) tracking and detecting a change in the mean of a sequence of observations. Here, the test-statistic (R_n) is a first-order autoregressive version of the observation process (Z_n) to be tracked with smoothing effected by an experimentally chosen parameter (λ) :

$$R_n = (1 - \lambda)R_{n-1} + \lambda Z_n, \quad n \geq 1$$

and $R_0 = 0$. The test-statistic is tested continuously against a threshold A_γ and change is declared at the first instant the test-statistic exceeds the threshold:

$$(4.2) \quad \tau_{\text{EWMA}} = \inf \{n \geq 1 : R_n \geq A_\gamma\}.$$

A_γ is chosen to ensure that the average run length (ARL) to a false alarm is at least γ . Small values of λ usually work best in changepoint detection as they smoothen small changes and enhance big changes (Srivastava and Wu, 1997).

Since we are interested in tracking a change in the latent attributes, we focus on an attribute that reflects the *resilience* of the group, and another that reflects the *level of coordination* in the group (Santos, 2011; Lindberg, 2010). In particular, the ability of a group to sustain terrorist activity over a number of days reflects its capacity to rejuvenate itself from asset (manpower, material, and skill-set) losses and to modify strategies and plans to counter establishment actions. On the other hand, the ability of the group to launch multiple attacks over a given time-period reflects its capacity to coordinate these assets necessary for simultaneous action often over a wide geographic area. That is, mathematically speaking, the focus is on: i) X_n , the number of days of terrorist activity, and ii) Y_n , the total number of attacks, both within the δ -day time-window Δ_n :

$$\begin{aligned} X_n &= \sum_{i \in \Delta_n} \mathbb{1}(\{M_i > 0\}) \\ Y_n &= \sum_{i \in \Delta_n} M_i \end{aligned}$$

for $n = 1, 2, \dots$ and where $\mathbb{1}(\cdot)$ denotes the indicator function of the set under consideration. Note that Y_n/δ is the average number of attacks per day and thus Y_n is a reflection of the intensity of attacks launched by the group. In general, X_n is more indicative of resilience in the group, whereas Y_n captures the level of coordination better.

Specializing the EWMA approach of (4.2) to spurt detection in the activity profile of a terrorist group with X_n and Y_n as observations, two parameters $\{\lambda_1, \lambda_2\} \in [0, 1]$ are chosen appropriately and used to update the variables $R_{1,n}$ and $R_{2,n}$ as follows:

$$\begin{aligned} R_{1,n} &= (1 - \lambda_1)R_{1,n-1} + \lambda_1 X_n \\ R_{2,n} &= (1 - \lambda_2)R_{2,n-1} + \lambda_2 Y_n \end{aligned}$$

for $n \geq 1$ with $R_{1,0} = 0 = R_{2,0}$. The best choices of smoothing parameters λ_1 and λ_2 for changepoint detection are obtained experimentally/numerically. We propose³ three stopping-times for declaring change: one based on $R_{1,n}$, another based on $R_{2,n}$, and the third on a weighted combination (with weights α and $\sqrt{1 - \alpha^2}$, $\alpha \in [0, 1]$) of the two test-statistics:

$$\begin{aligned} \tau_1 &= \inf \{n \geq 1 : R_{1,n} \geq A_1\} \\ \tau_2 &= \inf \{n \geq 1 : R_{2,n} \geq A_2\} \\ \tau_{\text{weighted}} &= \inf \left\{ n \geq 1 : \alpha R_{1,n} + \sqrt{1 - \alpha^2} R_{2,n} \geq A \right\}, \end{aligned}$$

where the thresholds A_1, A_2 , and A are chosen to meet the corresponding ARL constraints. While design of the threshold(s) requires further study, experimental studies suggest that a threshold of the form

$$\{A_1, A_2, A\} = \mathcal{O}(\log(\gamma))$$

ensures that $\{\text{ARL}(\tau_1), \text{ARL}(\tau_2), \text{ARL}(\tau_{\text{weighted}})\} = \mathcal{O}(\gamma)$. Since τ_{weighted} combines the information contained in both $\{X_n\}$ and $\{Y_n\}$, it should empirically be a better test than both τ_1 and τ_2 . Nevertheless, all the three tests could be of potential utility depending on the nature of the terrorist group.

4.2. Model-Based Strategy. We are now interested in model-based approaches for spurt detection. For this, note that the rare nature of terrorism suggests that most terrorist groups tend to be in an *Inactive* state for far longer than in an *Active* state. Thus, it is reasonable⁴ to assume that $\mathbf{S}_i = 0$ for long stretches of time and $\gamma_0 \approx 0$ where

$$(4.3) \quad \gamma_j \triangleq \text{P}(M_i > 0 | \mathbf{S}_i = j), \quad j = 0, 1.$$

Over such a long stretch where $\mathbf{S}_i = 0$, an elementary consequence of (4.3) is that X_n is a binomial random variable with parameters δ and γ_0 :

$$\text{P}(X_n = k) = \binom{\delta}{k} \cdot (\gamma_0)^k \cdot (1 - \gamma_0)^{\delta - k}.$$

³Note that other combinations such as the max statistic, min statistic, etc., can also be considered. We will not get into these details here.

⁴Note that in (4.3), we have not made any specific assumptions on the distribution of M_i . In fact, we have only labeled the quantity in (4.3) as γ_j .

If δ is sufficiently large (typical values used in subsequent case-studies are $\delta = 15$ to 30 days) so that the binning/edge effects can be neglected, X_n can be well-approximated by a Poisson random variable with rate parameter $\delta\gamma_0$. In fact, we have the following bound (Teerapabolarn, 2012, Cor. 3.2) on the approximability of the binomial distribution by Poisson for $k = 0, \dots, \delta$:

$$(4.4) \quad |\mathbb{P}_{\text{Bin}}(X_n = k) - \mathbb{P}_{\text{Poi}}(X_n = k)| \leq \min \left(1 - e^{-\delta\gamma_0}, \frac{\delta\gamma_0}{k} \right) \cdot \gamma_0 \approx 0.$$

Equivalently, let $T_k, k = 1, 2, \dots$ denote the time to the k th day of terrorist activity (with T_0 set to $T_0 = 0$). Then, $\Delta T_k = T_k - T_{k-1}$ denotes the number of days of wait to the subsequent day of terrorist activity (inter-arrival duration) and ΔT_k is approximately exponential with mean $1/\gamma_0$. While a similar reasoning suggests that ΔT_k in the *Active* state is exponential with mean $1/\gamma_1$, this fit is bound to be good only in the first-order sense because a terrorist group is expected to stay in the *Active* state for relatively shorter durations and $\gamma_1 \gg \gamma_0$. Rephrasing the above conclusions, a discrete-time Poisson process model is a good model for the days of terrorist activity, especially in the *Inactive* state.

Under the same assumptions (as above), in the *Inactive* state, Y_n can be rewritten as

$$Y_n = \sum_{i \in \mathcal{A}_n} M_i$$

where $\mathcal{A}_n \subset \Delta_n$ is the set of days of activity in Δ_n with $|\mathcal{A}_n| = X_n$. Thus, Y_n can be approximated as a compound Poisson random variable (Cox and Isham, 1980) whose density can be expressed in terms of the density function of M_i . For example, if M_i is independent and identically distributed (i.i.d.) as geometric with $\mathbb{P}(M_i = k | \mathbf{S}_i = 0) = (1 - \gamma_0) \cdot (\gamma_0)^k, k \geq 0$, a straightforward computation (see Supplementary Part A) shows that

$$\mathbb{P}(Y_n = r) = (1 - \gamma_0)^\delta \cdot (\gamma_0)^r \cdot \binom{\delta + r - 1}{r}, \quad r \geq 0.$$

Similarly, the joint density of (X_n, Y_n) can be written as

$$\mathbb{P}(X_n = k, Y_n = r) = (1 - \gamma_0)^\delta (\gamma_0)^r \cdot \binom{\delta}{k} \cdot \binom{r - 1}{r - k}, \quad r \geq k,$$

where the condition $r \geq k$ ensures that at least one attack occurs on a day of activity. With the more general hurdle-based geometric model, where

$$\mathbb{P}(M_i = k | \mathbf{S}_i = 0) = \begin{cases} 1 - \gamma_0, & k = 0 \\ \gamma_0(1 - \mu_0) \cdot (\mu_0)^{k-1}, & k \geq 1, \end{cases}$$

the joint density is given as

$$\mathbb{P}(X_n = k, Y_n = r) = \binom{\delta}{k} \binom{r - 1}{r - k} \cdot (1 - \gamma_0)^{\delta - k} (\gamma_0)^k \cdot (1 - \mu_0)^k (\mu_0)^{r - k}, \quad r \geq k.$$

Replacement of γ_0 and μ_0 with γ_1 and μ_1 in the *Active* state works subject to the same issues/constraints as stated earlier.

While more detailed changepoint detection strategies such as the Cumulative Sum (CUSUM) or the Shiryaev-Roberts (SR) test (Tartakovsky and Veeravalli, 2005) can be constructed by using the statistics of X_n and Y_n , we now propose an alternate approach to estimate the hidden states with the observations being $\{X_n\}$, $\{Y_n\}$, and the joint sequence $\{(X_n, Y_n)\}$ under different modeling assumptions on $\{M_i\}$. For this, we first apply the classical HMM formulation (Rabiner, 1989) where the Baum-Welch algorithm is used to learn the parameters that determine the density function of the observations. For the Baum-Welch algorithm to converge to a local maximum (with respect to the log-likelihood function) in the parameter space, a sufficient condition is that the density function of the observation be log-concave (Rabiner, 1989, Sec. IVA, p. 267). In Supplementary Part A, under the i.i.d. geometric and hurdle-based geometric models for $\{M_i\}$, it is established that all the three density functions are log-concave and hence also robust to initial estimates of the parameters. Further, an iterative update for the parameter estimates is also established under these two models. With the converged Baum-Welch parameter estimates as the initialization, the Viterbi algorithm is then used to estimate the most probable state sequence given the observations. The output of the Viterbi algorithm is a state estimate for the period of interest

$$\left\{ \mathbf{S}_i = \hat{s}_n \in \{0, 1\} \text{ for all } i \in \Delta_n \text{ and } n = 1, \dots, K \right\}.$$

A state estimate of 1 indicates that the group is *Active* over the period of interest whereas an estimate of 0 indicates that the group is *Inactive*. Transition between states indicates spurt/downfall in the activity.

5. Case-Studies. We now undertake two case-studies on the fit of a discrete-time Poisson point process model for the days of terrorist activity. For this, we classify terrorist activities as reported in the RDWTI and GTD that catalogue activities by different groups across the world (RDWTI; LaFree and Dugan, 2007).

5.1. FARC. We start with FARC, one of the most-feared and adaptive terrorist groups in the Americas. FARC is a Marxist-Leninist (pro-leftist) terrorist group active over a large geographical spread in Colombia and its neighborhood. FARC has a clearly laid out ideology that motivates and enables anti-government (and in general, anti-establishment) activities. We study the activities of FARC over the nine-year time-period from 1998 to 2006. This period covers a total of 604 terrorist incidents in the RDWTI with a yearly break-down of 44 incidents in 1998, 18 in 1999, 45 in 2000, 27 in 2001, 217 in 2002, 57 in 2003, 33 in 2004, 66 in 2005, and 97 in 2006, respectively. The reasons why FARC and the 1998 to 2006 time-period have been chosen for our study are provided in Supplementary Part B.

As explained in Section 4, the activity profile of FARC can be modeled as a discrete-time sampled Poisson point process. In Fig. 3, we test the fit of this model by studying: i) the Quantile-Quantile (Q-Q) plots comparing the sample inter-arrival duration between consec-

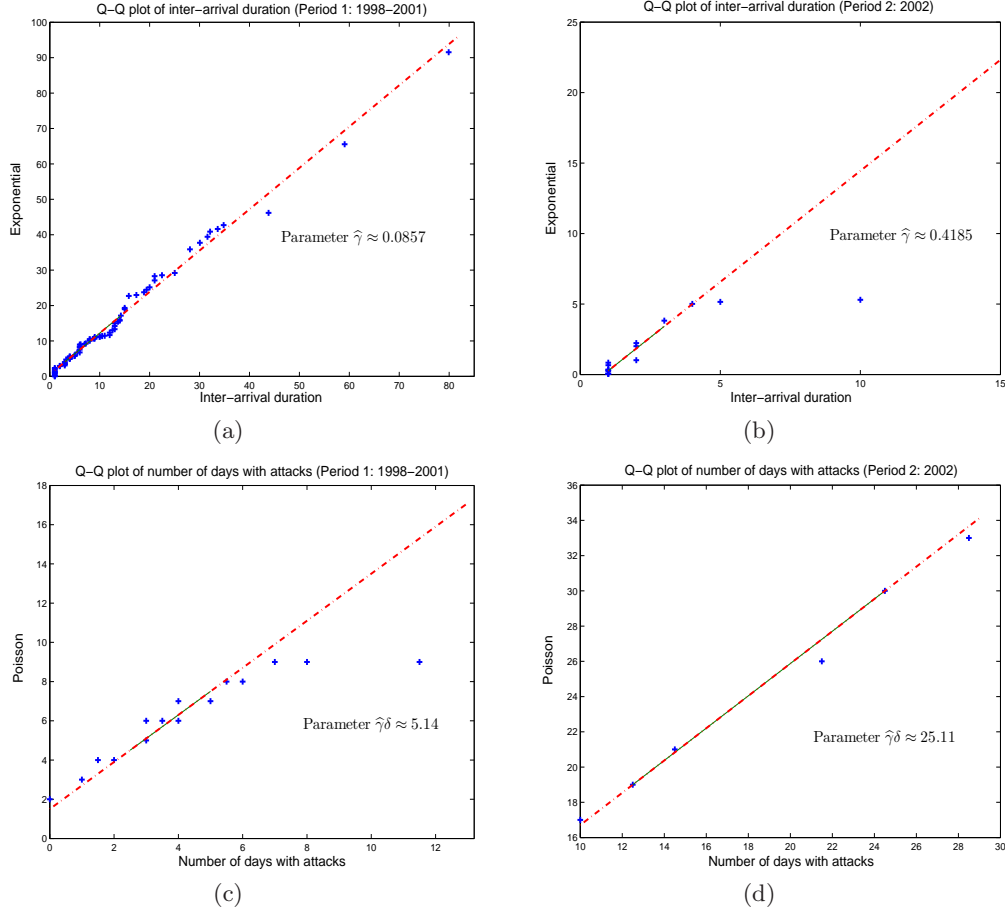


FIG 3. Q-Q plots of duration between days of terrorist activity with respect to a theoretical exponential random variable for the period from: (a) 1998 to 2001 corresponding to normal terrorist activity, and (b) 2002 corresponding to a spurt in terrorist activity. Q-Q plots of number of days of terrorist activity over a $\delta = 60$ day time-window with respect to a theoretical Poisson random variable for the same two scenarios are presented in (c) and (d).

utive days of terrorist activity and an exponential random variable with an appropriately-defined rate parameter ($\hat{\gamma}$), and ii) the Q-Q plots comparing the number of days of terrorist activity over a time-window of δ days and a Poisson random variable with parameter $\delta\hat{\gamma}$. For both sets of Q-Q plots, we consider two scenarios: the first (the 1998 to 2001 period) corresponding to a period of “normal” terrorist activity, and the second (the 2002 period) corresponding to a spurt in terrorist activity.

Figs. 3(a) and (b) compare the Q-Q plots of the sample inter-arrival duration under these two scenarios with an exponential random variable. The rate parameter used for the exponential is

$$(5.1) \quad \hat{\gamma} = \frac{1}{\mathbb{E}[\Delta T_k]}$$

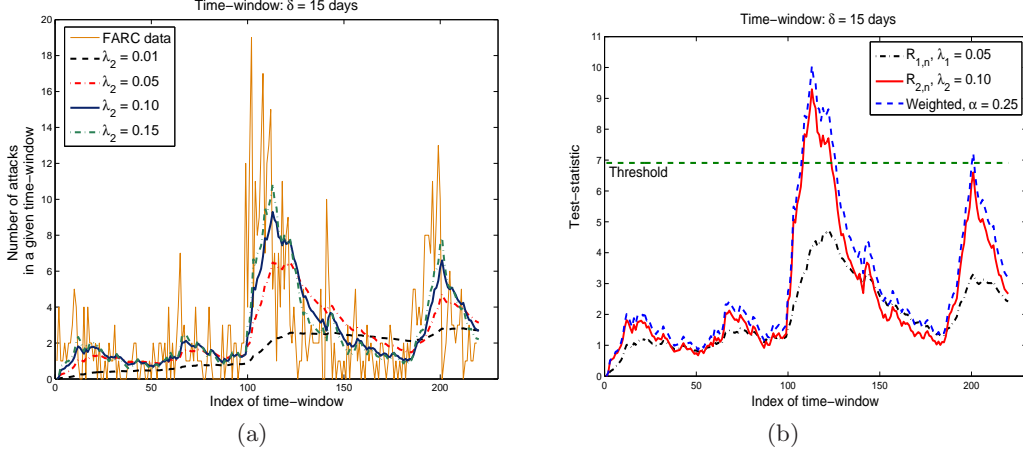


FIG 4. (a) Tracking properties of $R_{2,n}$ in response to changes in Y_n as a function of the smoothing parameter, (b) Performance of the three EWMA tests in spurt detection for FARC data ($\delta = 15$ days is used).

where $E[\Delta T_k]$ is the sample mean of the inter-arrival duration over the considered period. From Fig. 3, we note that both in periods of normal as well as a spurt in activity, the Q-Q plot is nearly linear with a few sample outliers in the tails. These outliers indicate that an exponential model for inter-arrival duration is not accurate because of the heavy tails. Nevertheless, to a first-order, an exponential random variable serves as a good fit for the inter-arrival durations. Our numerical study leads to the following estimates for γ_\bullet : (a) $\hat{\gamma} \approx 0.0857$, (b) $\hat{\gamma} \approx 0.4185$, suggesting that a spurt in activity is associated with an increase in the rate parameter. Similarly, in Figs. 3(c) and (d), we compare the Q-Q plots of the number of days of terrorist activity over a $\delta = 60$ day time-window under the same scenarios (as above) with a theoretical Poisson random variable of parameter $\delta \hat{\gamma} = 60 \hat{\gamma}$. While we observe some outliers in the tail quantiles, a Poisson random variable seems to be a good first-order fit for the number of days of terrorist activity.

We now study the performance of the proposed EWMA approach with the FARC data. In the first study, presented in Fig. 4(a), we study how the tracking properties of $R_{2,n}$ (to changes in Y_n) behave with different values of the smoothing parameter λ_2 . The true behavior of Y_n is also provided as a benchmark to compare the behaviors of $R_{2,n}$. In general, small values of λ_2 smoothen big changes in Y_n considerably, whereas large values of λ_2 ensure that the evolution of $R_{2,n}$ is vulnerable to even small changes in Y_n . As can be seen from the plot, a value of $\lambda_2 = 0.10$ trades-off the phenomena of smoothing out of small changes and enhancing of big changes, and thus renders the test-statistic $R_{2,n}$ effective from a changepoint detection perspective.

In Fig. 4(b), we study the performance of the three EWMA tests proposed in this work. We plot the test-statistics: $R_{1,n}$ with $\lambda_1 = 0.05$ for τ_1 , $R_{2,n}$ with $\lambda_2 = 0.10$ for τ_2 , and $\alpha R_{1,n} + \sqrt{1 - \alpha^2} R_{2,n}$ with $\alpha = 0.25$ for τ_{weighted} . The threshold is designed as $\{A_1, A_2, A\} = 3 \log(\gamma)$ for $\gamma = 10$. From Fig. 4(b), we see that an appropriate weighted combination of the metric that captures resilience and the level of coordination in the group performs

better than either test-statistic taken separately (with the same threshold for all the three tests). While with the FARC data, the weighted sum performs only marginally better than the resilience-based metric, in general, we expect τ_{weighted} to significantly improve the performance over either τ_1 or τ_2 . But more importantly, the EWMA approach detects only *persistent* changes or changes that last for a sufficiently long duration so that the changepoint detection approach can work accurately. In other words, the major spurts in FARC activity are detected, while the approach *cannot* detect the minor spurts. This is because the approach does not incorporate the statistical information of $\{X_n\}$ or $\{Y_n\}$.

TABLE 1

State classification of FARC with the geometric and hurdle-based geometric models and $\{(X_n, Y_n)\}$ as observations for different time-window periods (δ)

δ (in days)	Parameters learned		Length of training-set (N time-windows)	No. of <i>Active</i> states (N_{spurt} time-windows)	Fractional activity ($f = \frac{N_{\text{spurt}} \cdot \delta}{N}$)
	$\hat{\gamma}_0$	$\hat{\gamma}_1$			
1	0.0924	0.3597	3286	483	0.1469
5	0.0919	0.3584	657	140	0.2130
7	0.0924	0.3598	469	89	0.1895
10	0.0911	0.3568	328	73	0.2221
15	0.0924	0.3605	219	46	0.2099
25	0.0908	0.3592	131	26	0.1977
30	0.0930	0.3593	109	20	0.1825

δ (in days)	Parameters learned				Length of training-set (N time-windows)	No. of <i>Active</i> states (N_{spurt} time-windows)	Fractional activity ($f = \frac{N_{\text{spurt}} \cdot \delta}{N}$)
	$\hat{\gamma}_0$	$\hat{\mu}_0$	$\hat{\gamma}_1$	$\hat{\mu}_1$			
1	0.0950	0.0752	0.3982	0.3083	3286	484	0.1472
5	0.0942	0.0730	0.3934	0.3066	657	140	0.2130
7	0.0949	0.0742	0.3956	0.3081	469	87	0.1853
10	0.0936	0.0724	0.3891	0.3044	328	71	0.2160
15	0.0958	0.0759	0.4019	0.3103	219	42	0.1917
25	0.0934	0.0738	0.3957	0.3046	131	26	0.1977
30	0.0954	0.0794	0.3966	0.3073	109	20	0.1825

To overcome this deficiency, as explained in Section 4.2, a geometric model is assumed for $\{M_i\}$ and the Baum-Welch algorithm is used to learn the underlying γ_j with $\{X_n\}$, $\{Y_n\}$, and $\{(X_n, Y_n)\}$ over a given δ -day time-window as training data. Specifically, Table 1 summarizes the parameter estimates for different δ values when $\{(X_n, Y_n)\}$ is used as a training-set. It is to be noted that the learned parameter values remain reasonably stable across a large range of δ (1 to 30 days) values. The performance of the Baum-Welch algorithm is also robust to both the length of the training-set as well as the initialization. Further, the parameter values also remain essentially independent of whether $\{X_n\}$, $\{Y_n\}$, or $\{(X_n, Y_n)\}$ is used for training. For example, with $\delta = 15$ and $\{(X_n, Y_n)\}$ as training data, the parameter values learned are $\hat{\gamma}_0 = 0.0924$ and $\hat{\gamma}_1 = 0.3605$, whereas with $\{X_n\}$, these values are $\hat{\gamma}_0 = 0.0941$ and $\hat{\gamma}_1 = 0.3861$. This observation is not entirely surprising since $\{M_i\}$ is assumed to come from a one parameter model family and conditioned on

one of the two variables (X_n or Y_n), the other variable adds no significant new information about the model parameter.

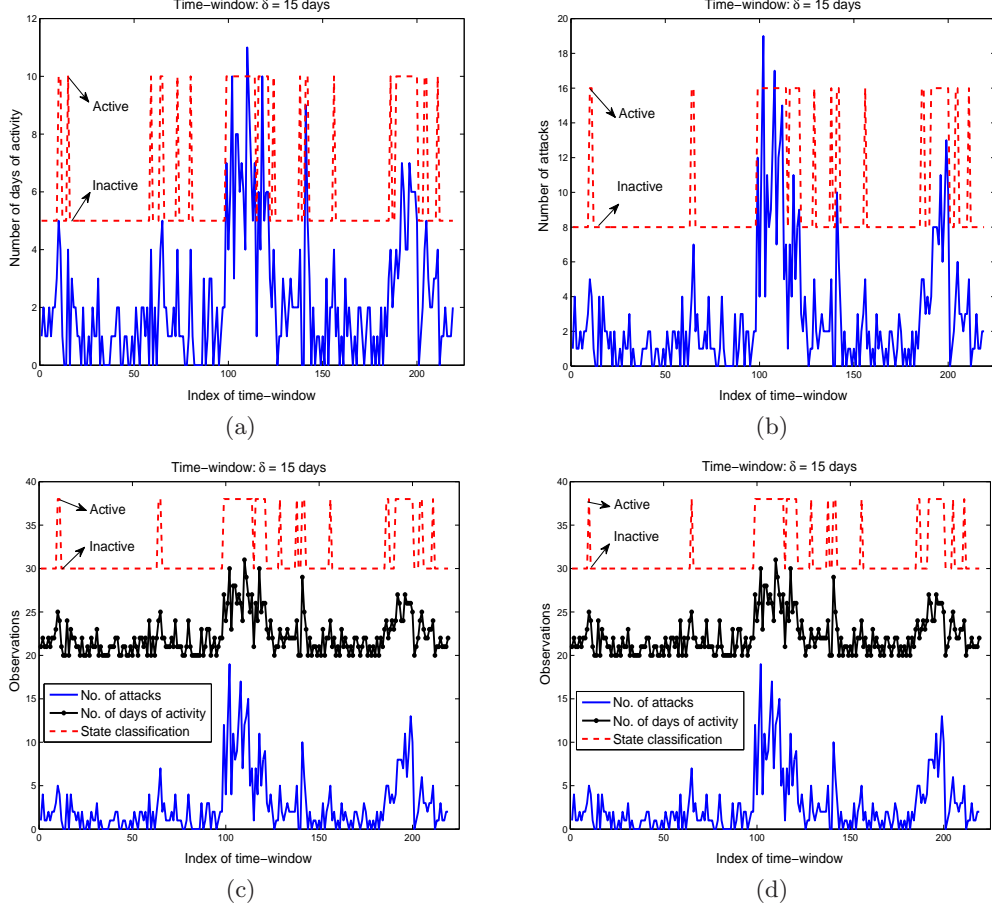


FIG 5. State classification via Viterbi algorithm with the geometric model for $\{M_i\}$; observations being a) $\{X_n\}$, b) $\{Y_n\}$, and c) $\{(X_n, Y_n)\}$ with parameters learned by Baum-Welch algorithm. d) State classification with the hurdle-based geometric model for the observation sequence $\{(X_n, Y_n)\}$.

The converged Baum-Welch parameter estimates are then used to initialize the Viterbi algorithm to obtain the most probable state sequence for FARC. Figs. 5(a)-(c) illustrate the state classification for $\delta = 15$ with $\{X_n\}$, $\{Y_n\}$, and $\{(X_n, Y_n)\}$ as observations. As can be seen from Fig. 5, in a clear departure from changepoint detection with the EWMA approach, the state classification approach detects even *small* and *non-persistent* changes. Further, the Viterbi algorithm declares 51 and 46 of the 219 time-windows as *Active* with $\{X_n\}$ and $\{Y_n\}$, respectively, whereas the joint sequence $\{(X_n, Y_n)\}$ results in the same classification as $\{Y_n\}$. Table 1 summarizes the number of time-windows classified as *Active* for different δ values with $\{(X_n, Y_n)\}$ as observations. This study also suggests that $\delta = 10$ to 15 with the HMM approach optimally trades-off the twin objectives of detecting minor spurts in the activity profile of FARC (larger N_{spurt}) as well as minimizing the number of

Active state classifications that require further attention (smaller fractional activity f).

We now consider a more general two-parameter hurdle-based geometric model for $\{M_i\}$ that potentially allows the joint sequence $\{(X_n, Y_n)\}$ to result in better inferencing on the states than either $\{X_n\}$ or $\{Y_n\}$. As before, the Baum-Welch algorithm is used to learn the underlying parameters with different values of δ and $\{(X_n, Y_n)\}$ as training data. Table 1 summarizes these parameter estimates and as was the case earlier, it can be seen that the estimates remain reasonably stable across δ . The parameter estimates are then used with the Viterbi algorithm to infer the most probable state sequence (see Fig. 5(d) for state classification in the $\delta = 15$ setting). While Fig. 5(d) and Table 1 show that the classification with the hurdle-based model agrees with the simpler geometric model for many δ values, in the setting of interest ($\delta = 7$ to 15 days), the hurdle-based model is more conservative in declaring an *Active* state. The four time-windows of disagreement between the hurdle-based geometric and geometric models for $\delta = 15$ correspond to the boundary of minor spurts (time-windows 11, 64, 191, and 204) where the hurdle-based model is more conservative in declaring an *Active* state, whereas the geometric model is *trigger-happy* and overcautious.

To provide further intuition on the classification behind the Viterbi algorithm, we now propose an *ad hoc* approach to classify states, the outcome of which mirrors the outcome of the Viterbi algorithm. In this approach, state classification is done in two steps. In the first step, $\{\Delta T_k\}$ are used as observations with the Viterbi algorithm to result in a state estimate for the days of activity

$$\left\{ \tilde{s}_k \in \{0, 1\} \quad \text{for all } k : M_k > 0 \right\}.$$

In the second step, the following re-estimation formula provides an estimate of the states over the δ -day time-window:

$$(5.2) \quad \bar{s}_n = \begin{cases} 1, & \text{if } \frac{\sum_{i \in \Delta_n} M_i \tilde{s}_i}{\sum_{i \in \Delta_n} M_i} > \eta \\ 0, & \text{if otherwise.} \end{cases}$$

Note that the above formula is well-defined as it needs \tilde{s}_i only for the days of activity since $M_i = 0$ otherwise. Further, η is a threshold that is chosen appropriately (typical values used in our study are $\eta = 0.5, 0.75$, etc).

Since the Viterbi algorithm in the first step uses only $\{\Delta T_k\}$ as observations, the frequency of attacks perpetrated by the group (X_n or the group's resilience) is the most important factor that determines the initial state assignment $\{\tilde{s}_k\}$. In fact, the number of attacks on a day (Y_n or the group's level of coordination) has no impact on $\{\tilde{s}_k\}$. To compensate for not using this information, the re-assignment procedure in (5.2) is biased to the *Active* state if more attacks happen in a time-window. For example, if $M_i = m \geq 1$ for all $i \in \Delta_n$, \bar{s}_n is declared 1 only if an η fraction of the days of activity in Δ_n are declared as *Active* in the first step. Similarly, if there exists a single day of heightened activity in Δ_n such that $M_i \gg M_j$ for some i and all $j \neq i$, \bar{s}_n is declared 1 provided $\tilde{s}_i = 1$ in the first step. In other words, the state declaration $\{\bar{s}_n\}$ combines both the facets that are responsible for a spurt in the activity profile: an increased frequency of attacks, or more attacks on a

day. With $\delta = 15$ for the FARC dataset, the *ad hoc* approach declares as *Active* all the 42 time-windows classified by the Viterbi algorithm as *Active*. In addition, it also classifies 32 other time-windows (that are declared by the Viterbi algorithm as *Inactive*) also as *Active*. Thus, while the proposed approach is *trigger-happy* relative to the Viterbi algorithm, it nevertheless provides a certain intuition on how the resilience and coordination aspects of the group could be combined in understanding the group dynamics. While this combination of the two facets of the group is *ad hoc*, the proposed approach opens up the exploration of a more systematic combination of these facets. This will be the subject of future work.




FIG 6. (a) Number of days of terrorist activity, and (b) Total number of attacks in a $\delta = 25$ day time-window as a function of time for Shining Path.

5.2. *Shining Path*. The second case-study is the activity profile of Sendero Luminoso (more popularly known by its English name of *Shining Path*), one of the major terrorist groups in Peru for the sixteen year time-period from 1981 to 1996. Shining Path was guided by the Maoist (pro-leftist) ideology with a strong inclination to attack targets that had been perceived as “interventionist” or “establishmentist.” Thus, it had been primarily involved in guerilla warfare on such diverse sets of targets as non-Peruvians as well as coordinated attacks on: i) government installations, ii) embassies and foreign missions in Peru, iii) cultural centers of, and aid organizations funded by foreign countries, iv) multi-national as well as foreign-owned companies, etc. Further, Shining Path was involved in the drug trafficking and trans-shipment business as a means to fund its activities ([World Drug Report, 2010](#)). With a focus on the sixteen-year period between 1981 and 1986, the RDWTI reports⁵ a total of 163 incidents with a yearly breakdown of 10 incidents in 1981, 7 in 1982, 10 in 1983,

⁵After much of this work had been completed, the authors learned that the GTD could provide a better dataset for Shining Path than RDWTI. In fact, the GTD reports at least 4000 attacks over the same time-period for Shining Path. Processing of this data is manually intensive and a future work on comparative lessons on different types of terrorist groups is planned to address this data. Thus, much of the ensuing study should be seen in the light of what is possible with the proposed algorithms rather than as a reflection of the quality of the dataset used in the study.




FIG 7. *Q-Q plots of duration between days of terrorist activity with respect to a theoretical exponential distribution for the period from: a) 1981 to 1986, b) 1987 to 1990, c) 1991, and d) 1992 to 1996 corresponding to different stages of evolution of Shining Path.*

7 in 1984, 3 in 1985, 12 in 1986, 19 in 1987, 7 in 1988, 18 in 1989, 10 in 1990, 31 in 1991, 10 in 1992, 14 in 1993, 2 in 1994, 2 in 1995, and 1 in 1996. The choice of Shining Path and the 1981 to 1996 time-period are motivated in Supplementary Part B.

In Fig. 6, we plot the number of days of terrorist activity and the total number of attacks over a $\delta = 25$ day time-window. In Fig. 7, we test the fit of inter-arrival duration between successive days of terrorism incidents with respect to an exponential random variable with parameter $\hat{\gamma}$ estimated as in (5.1). As noted in Supplementary Part B, the evolution of Shining Path can be partitioned into four distinct phases. The inter-arrival duration in each phase can be distinctly modeled as an exponential random variable and Fig. 7 shows that this partitioning is reasonable. As in the case with FARC, the exponential random variable is a good first approximation as the tails are not well-modeled with this random variable. Our numerical study leads to the following estimates for γ_{\bullet} in the four phases: a) $\hat{\gamma} \approx 0.0211$, b) $\hat{\gamma} \approx 0.0344$, c) $\hat{\gamma} \approx 0.0755$, d) $\hat{\gamma} \approx 0.0158$.

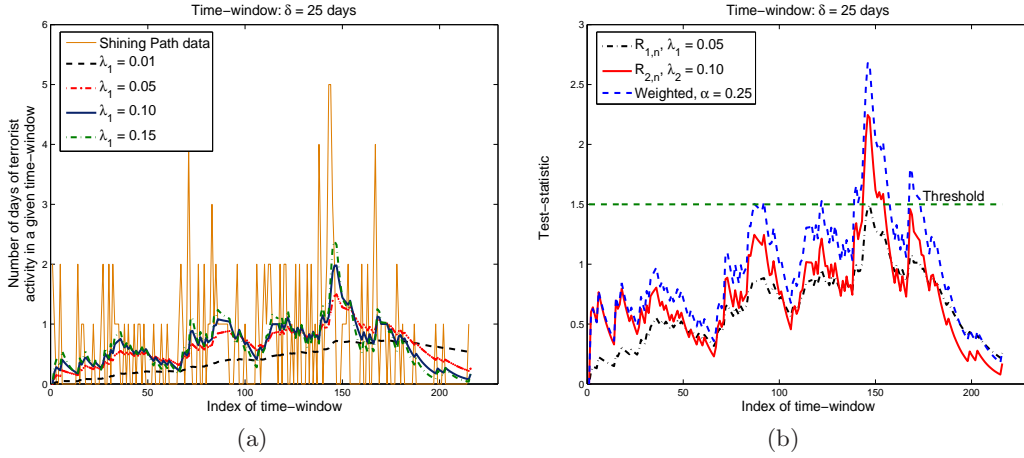



FIG 8. (a) Tracking properties of $R_{1,n}$ in response to changes in X_n as a function of the smoothing parameter, (b) Performance of the three EWMA tests in spurt detection for Shining Path data ($\delta = 25$ days).

As elucidated with the FARC dataset, the data corresponding to Shining Path is studied in the following experiments. In Fig. 8(a), the tracking properties of $R_{1,n}$ for different values of λ_1 are presented along with the number of days of terrorist activity over a $\delta = 25$ day time-window. From this figure, we see that a small value of λ_1 such as $\lambda_1 = 0.05$ or 0.10 is best suited in changepoint detection. In Fig. 8(b), the test-statistic corresponding to τ_1 , τ_2 , and τ_{weighted} are presented along with a threshold that is chosen appropriately. Using the HMM approach with the hurdle-based geometric model described in Section 4.2, the Baum-Welch algorithm results in parameter estimates as in Table 2 for different values of δ . State classification via the Viterbi algorithm using these estimates results in an *Active/Inactive* classification for Shining Path, e.g., Fig. 9 displays a typical classification for $\delta = 25$. It is important to note that while four distinct phases are identified in the evolution of Shining Path, only a $d = 2$ -state HMM is studied here. For the purpose of spurt/downfall detection, even this coarse and parsimonious model is sufficient. As can be seen with FARC data, even small and non-persistent changes are detected by the HMM approach, further confirming its usefulness.

TABLE 2
State classification of Shining Path with the hurdle-based geometric model and $\{(X_n, Y_n)\}$ as observations for different time-window periods (δ)

δ (in days)	Parameters learned				Length of training-set (N time-windows)	No. of <i>Active</i> states (N_{spurt} time-windows)	Fractional activity ($f = \frac{N_{\text{spurt}} \cdot \delta}{N}$)
	$\hat{\gamma}_0$	$\hat{\mu}_0$	$\hat{\gamma}_1$	$\hat{\mu}_1$			
20	0.0262	0.1026	0.1000	0.6667	268	69	0.2575
25	0.0264	0.1019	0.0800	0.6667	215	35	0.1628
28	0.0264	0.1019	0.0714	0.6667	192	15	0.0781
30	0.0262	0.1026	0.0667	0.6667	179	14	0.0782
32	0.0202	0.0703	0.1072	0.2255	168	38	0.2262

FIG 9. State classification with Shining Path data ($\delta = 25$ days).

6. Revisiting Some of the Modeling Issues. We now re-explore some of the modeling issues that shed further light on development of good models for the activity profile of terrorist groups. In this section, the primary focus is on the FARC dataset.

6.1. Clustering of Attacks. A central premise in many work on terrorism modeling is that the current activity of a terrorist group is influenced by its past activity. One consequence of this premise is that the attacks perpetrated by the group are *clustered* (Midlarsky, 1978; Midlarsky, Crenshaw and Yoshida, 1980). Ripley's $K(\cdot)$ function is a statistical tool for measuring the degree of clustering (aggregatedness) or inhibition (regularity) in a point process as a function of inter-point distance (Diggle, 2003). Specifically, if λ is the intensity of the point process, $\lambda K(h)$ is the expected number of other points within a distance h of a randomly chosen point of the process:

$$K(h) \triangleq \frac{1}{\lambda} \cdot \mathbb{E}[\text{Number of other points within distance } h \text{ of a randomly chosen point}].$$

Expressions for $K(h)$ can be derived for a number of stationary point process models (Dixon, 2002). For example, in the case of a one-dimensional/temporal point process that is completely random (where points are distributed uniformly and independently in time), it can be shown that $K(h) = 2h$. A two-dimensional complete random spatial point process leads to $K(h) = \pi h^2$.

In the context of an activity profile, $K(h)$ can be estimated as

$$\begin{aligned} \hat{K}(h) &= \frac{1}{\hat{\lambda} \cdot N} \sum_i \sum_{j \neq i} \mathbb{1}(|t_i - t_j| \leq h) \\ &= \frac{1}{(\hat{\lambda})^2 \cdot \mathcal{N}} \sum_i \sum_{j \neq i} \mathbb{1}(|t_i - t_j| \leq h) \end{aligned} \quad (6.1)$$

where t_i is the i th day of activity, N is the number of attacks in the observation period of \mathcal{N} days (that is, $[t_1, \dots, t_N] = [\text{Day } 1, \dots, \text{Day } \mathcal{N}]$), and $\hat{\lambda} = \frac{N}{\mathcal{N}}$ is an estimate of the intensity

(rate) of the process. Significant deviations of $\hat{K}(h)$ from $2h$ indicate that the hypothesis of complete randomness becomes untenable with observed data and more confidence is reposed on clustering (if $\hat{K}(h) \gg 2h$) or inhibition (if $\hat{K}(h) \ll 2h$).

While the definition in (6.1) assumes a homogenous point process, extensions to an inhomogenous point process have also been proposed (Baddeley, Moller and Waagepetersen, 2000; Veen and Schoenberg, 2006), where the point process is re-weighted by the reciprocal of the non-constant intensity function to offset the inhomogeneity. Further, to compensate for edge effects due to points outside the observation period being left out in the numerator of (6.1), various edge-correction estimators have also been proposed in the literature. Combining these two facets, we have the following estimator⁶ for $K(h)$:

$$(6.2) \quad \hat{K}(h) = \frac{1}{\mathcal{N}} \sum_i \sum_{j \neq i} \frac{\mathbb{1}(|t_i - t_j| \leq h)}{\hat{p}_i \hat{p}_j w_{ij}}$$

where \hat{p}_i is the estimated probability of at least one attack on t_i . In this work, we use an edge-correction factor w_{ij} due to Ripley (Cressie, 1991, pp. 616-618) which reflects the proportion of the period centered at t_i and covering the j th day of activity that is included in the observation period:

$$w_{ij} = \begin{cases} 1, & \text{if } t_1 + R \leq t_i \leq t_N - R \\ \frac{t_N - t_i + R}{2R}, & \text{if } t_i > \max(t_N - R, t_1 + R) \\ \frac{R + t_i - t_1}{2R}, & \text{if } t_i < \min(t_N - R, t_1 + R) \\ \frac{t_N - t_1}{2R}, & \text{if } t_N - R \leq t_i \leq t_1 + R \end{cases}$$

where $R = |t_i - t_j|$.

In Fig. 10(a), we plot $K(h) - 2h$ computed as in (6.2) for the FARC dataset (with and without edge-correction) as a function of the inter-point distance h . In line with the observation in Porter and White for the Indonesia/Timor-Leste dataset, the plot here indicates that the FARC data is also clustered since $K(h) - 2h \gg 0$, thus motivating the self-exciting hurdle model of Porter and White. The main theme of this work, however, is that clustering is essentially a reflection of state transitions and the activity sub-profile (sub-series) of the group conditioned on a given state value is not clustered. To test this hypothesis, we study the behavior of $K(h)$ for the sub-series from the FARC dataset corresponding to the *Active* and *Inactive* states as classified by the methodology of Section 5.1.

Since the latent states of the group transition back and forth, an *Active* sub-series is constructed by patching together the group's activity profile in the *Active* state with the *jump* random variable between any two disjoint pieces of the activity profile modeled as Poisson with parameter $\hat{\lambda} = \frac{\mathcal{N}}{\mathcal{N}_{\text{act}}}$. For this sub-series, $K(h)$ is estimated using a formula analogous to (6.2), where \hat{p}_i and w_{ij} are re-estimated for the sub-series, and \mathcal{N} is replaced with \mathcal{N}_{act} :

$$\mathcal{N}_{\text{act}} = \min \left(\text{Number of days in } \textit{Active} \text{ state, Length of } \textit{Active} \text{ sub-series} \right).$$

⁶Note that Porter and White propose a one-sided estimator for $K(h)$ and they compare $\hat{K}(h)$ with h (instead of $2h$) to test for clustering/inhibition.




FIG 10. (a) $K(h)$ as a function of inter-point distance h for the activity profile of FARC and under the two states of the HMM, (b) Fractional change in state classification relative to addition of missing data from GTD.

A similar estimation process yields $K(h)$ for the *Inactive* sub-series. In contrast to the trends for the activity profile, the sub-series' of the activity profile indicate an inhibitory behavior (mild inhibition for the *Active* state and strong inhibition for the *Inactive* state). These conclusions should not be entirely surprising since very few attacks happen in the *Inactive* state. Thus, Fig. 10(a) prompts that the *classical* assumption of clustering of attacks needs to be revisited and the *causality* assumptions behind terrorism models studied more carefully across many groups/datasets.

6.2. Model for Observation Densities. We now develop simple models for the observation density under the two states. For this, we study the goodness-of-fit captured by the AIC for several models with support on the non-negative integers to describe data from FARC. In Table 3, we present the histogram of the number of days with ℓ ($\ell = 0, 1, \dots$) attacks per day for FARC data. Applying the state classification algorithm described in Section 4.2 with $\delta = 15$, FARC stays in the *Inactive* state for 2657 days and in the *Active* state for 630 days. Also, presented in the same table are the (rounded-off) expected number of days with ℓ attacks for the six models described in Supplementary Part A, along with the AIC for these model-fits. Table 4 presents the corresponding histograms for $\delta = 10$, with which FARC stays in the *Inactive* state for 2577 days and in the *Active* state for 710 days.

The ML parameter estimates for all the six models remain reasonably robust as δ is varied, which is in conformity with the stability of the converged Baum-Welch parameter estimates with δ (see Tables 1 and 2). Further, from Tables 3 and 4, in the *Inactive* state, it is seen that all the models except the shifted Zipf result in comparable fits. Specifically, the hurdle-based geometric model differs from the observed histogram in only one day and results in the second lowest AIC value. On the other hand, in the *Active* state, the hurdle-based geometric model produces the best fit with only the Pòlya model resulting

TABLE 3
Histogram of observed number of attacks per day for FARC data with different model-fits, $\delta = 15$ days

No. attacks (Inactive State)	Obs.	Poisson	Shifted Zipf	Geomet.	Pòlya	Hurdle- Based Zipf	Hurdle- Based Geomet.
0	2420	2421	2470	2430	2421	2420	2421
1	227	225	144	207	225	229	226
2	9	11	27	18	11	7	10
3	1	0	8	2	0	1	0
4	0	0	4	0	0	0	0
> 4	0	0	4	0	0	0	0
AIC		1690.34	1772.81	1696.74	1692.32	1692.58	1691.86
Parameter Estimate		0.0933	4.105	0.0854	$\hat{r}_0 = 24.4749,$ $\hat{y}_0 = 0.0038$	$\hat{\gamma}_0 = 0.0892,$ $\hat{y}_0 = 5.10$	$\hat{\mu}_0 = 0.0444,$ $\hat{\gamma}_0 = 0.0892$
No. attacks (Active State)	Obs.	Poisson	Shifted Zipf	Geomet.	Pòlya	Hurdle- Based Zipf	Hurdle- Based Geomet.
0	384	359	455	404	389	384	384
1	174	202	87	144	160	189	171
2	46	57	33	52	56	31	52
3	19	11	16	19	17	11	16
4	4	1	9	7	6	5	5
> 4	3	0	30	4	2	10	2
AIC		1313.88	1416.88	1291.73	1288.85	1308.09	1287.11
Parameter Estimate		0.5651	2.40	0.3611	$\hat{r}_1 = 1.4834,$ $\hat{y}_1 = 0.2759$	$\hat{\gamma}_1 = 0.3905,$ $\hat{y}_1 = 2.61$	$\hat{\mu}_1 = 0.3090,$ $\hat{\gamma}_1 = 0.3905$

in a comparable fit. In this setting, the one-parameter models over-estimate either the tail or the days of no activity, while the hurdle-based Zipf produces a heavier tail than what the data exhibits. In fact, the poorest fit in either state is obtained with the shifted Zipf suggesting that a heavy tail may not always be necessary. In contrast, the Indonesia/Timor-Leste data studied in (Porter and White, 2012) exhibits several extreme values (e.g., days with 36, 11, and 10 attacks) and the authors observe that a self-exciting hurdle-based Zipf model captures the heavy tails much better than other models. The FARC dataset used here shows a maximum of 7 attacks per day, whereas the Shining Path⁷ dataset shows a maximum of 3 attacks per day. Even simple non-self-exciting models are enough to capture these datasets sufficiently accurately.

This study suggests the following:

- If parsimony of the model is of critical importance, the geometric distribution serves as the best one-parameter model with the Poisson/shifted Zipf models either under-estimating or over-estimating the number of days with no activity in the *Active* state.
- If parsimony is not a critical issue and the data does not have (or has very few) extreme values, the hurdle-based geometric model serves as the best/near-best model in either state.

⁷Note that the data from GTD on Shining Path, which is more richer, is not expected to conform to the trend from RDWTI.

TABLE 4

Histogram of observed number of attacks per day for FARC data with different model-fits, $\delta = 10$ days

No. attacks (Inactive State)	Obs.	Poisson	Shifted Zipf	Geomet.	Pòlya	Hurdle- Based Zipf	Hurdle- Based Geomet.
0	2379	2376	2420	2385	2380	2379	2379
1	188	193	124	178	188	191	188
2	10	8	22	13	9	6	9
3	0	0	6	1	0	1	1
4	0	0	2	0	0	0	0
> 4	0	0	3	0	0	0	0
AIC		1478.87	1537.97	1481.36	1480.56	1484.02	1480.77
Parameter Estimate		0.0807	4.29	0.0747	$\hat{r}_0 = 5.0975$, $\hat{y}_0 = 0.0156$	$\hat{\gamma}_0 = 0.0768$, $\hat{y}_0 = 4.915$	$\hat{\mu}_0 = 0.0481$, $\hat{\gamma}_0 = 0.0768$
No. attacks (Active State)	Obs.	Poisson	Shifted Zipf	Geomet.	Pòlya	Hurdle- Based Zipf	Hurdle- Based Geomet.
0	425	407	513	456	437	425	425
1	213	226	97	163	187	226	206
2	45	63	37	58	61	34	57
3	20	12	18	21	18	11	16
4	4	2	11	8	5	5	4
> 4	3	0	34	4	2	9	2
AIC		1452.09	1599.89	1444.86	1435.32	1443.56	1430.33
Parameter Estimate		0.5577	2.3950	0.3580	$\hat{r}_1 = 1.8427$, $\hat{y}_1 = 0.2323$	$\hat{\gamma}_1 = 0.4014$, $\hat{y}_1 = 2.735$	$\hat{\mu}_1 = 0.2803$, $\hat{\gamma}_1 = 0.4014$

- However, if the data has several extreme values (as seen in [Porter and White](#)), the self-exciting hurdle model offers the best model-fit albeit at the expense of learning several model parameters.

6.3. *Inter-Arrival Duration.* We now study the efficacy of the HMM framework by testing the goodness-of-fit of the theoretical exponential random variables with respect to the inter-arrival duration between days of terrorist activity in either state. To avoid estimating the rate parameters of the exponential random variables from the data (which complicates the hypothesis tests), we use the fact from ([Seshadri, Csorgo and Stephens, 1969](#)) that if y_1, \dots, y_m are i.i.d. exponential random variables (with a given rate parameter), then

$$(6.3) \quad z_j = \frac{\sum_{i=1}^j y_i}{\sum_{i=1}^m y_i}, \quad j = 1, \dots, m$$

are i.i.d. uniformly distributed in $[0, 1]$. We then use a Kolmogorov-Smirnov (KS) test to study the fit between the empirical cumulative distribution function (CDF) of z_j and the uniform CDF ([Durbin, 1973](#)).

The KS test-statistic (denoted as KS_\bullet) and the critical value for the test (denoted as K_α) corresponding to a significance level α and computed using the standard asymptotic

formula are given as

$$KS_{\bullet} = \max_x \left| \frac{1}{n_{\bullet}} \sum_{i=1}^{n_{\bullet}} \mathbb{1}(z_i^{\bullet} \leq x) - x \right|$$

$$K_{\alpha} = \sqrt{\frac{-\frac{1}{2} \log\left(\frac{\alpha}{2}\right)}{n_{\bullet}}}$$

where n_{\bullet} is the number of samples and $\{z_i^{\bullet}\}$ are the transformed samples computed using (6.3) in either state. The results of applying the KS test to the two states are presented in Table 5. From this table, it is clear that the samples in either state fit the theoretical exponential assumption very accurately with the exponential model rejected in the *Active* (and *Inactive*) state(s) if the significance level α exceeds $\alpha > 0.3763$ (and $\alpha > 0.7905$), respectively.

TABLE 5
KS test-statistic in the two states for FARC dataset

	<i>Active</i>	<i>Inactive</i>
Total number of samples	245	236
No. of samples with $\Delta T_k \leq \beta$	$n_A = 234, \beta = 15$	$n_I = 192, \beta = 20$
KS statistic	0.0597	0.0492
<i>p</i> -value	0.3763	0.7905

On the other hand, by studying terrorism data for FARC over the time-period 1998 to 2005 from the MIPT database (a database comparable with RDWTI, but currently unavailable), Clauset and Gleditsch hypothesize that the inter-arrival duration between successive days of attacks “decreases consistently, albeit stochastically” with the cumulative number of events FARC has carried out – a measure of the group’s experience. Our initial studies indicate that while Clauset and Gleditsch’s hypothesis holds empirically true for $k \leq 25$ attacks that encompass the period Jan. 1 to June 27, 1998, it consistently increases through the subsequent period lasting till Mar. 10, 2000. Note that this is the precise time-period of increased U.S. funding to combat FARC and the drug economy through *Plan Colombia* (Haugaard, Isacson and Olson, 2005) and a mean increase in the time to the next day of activity indicates an impact of counter-terrorism efforts. The following period through Aug. 13, 2004 indicates a reversed trend of consistent decrease suggesting that the organizational dynamics of FARC had “adjusted” to the new reality of combat with the establishment. As seen earlier, such distinct changes in the organizational dynamics (associated with spurts and downfalls) are quickly identified by the approaches proposed in this work.

6.4. *Robustness of Proposed Approach to Missing Data.* We now study the robustness of the proposed spurt detection approach in terms of state classification to attacks that are not available in the database. Toward this goal, we treat the FARC data from RDWTI over the 1998 to 2006 period as the baseline dataset and add one missing day of activity per year from the GTD in a *sequential* manner and revisit state classification with the enhanced

dataset. Specifically, the fraction of missing data added in the j th (sequential) step is the ratio of the difference between new and old attacks and the baseline, and is defined as

$$\text{Frac. Missing Data}(j) \triangleq \frac{\sum_{i=1}^{\mathcal{N}} \widetilde{M}_i^j - M_i}{\sum_{i=1}^{\mathcal{N}} M_i}$$

where \widetilde{M}_i^j is the number of attacks on the i th day with data addition. Applying the state estimation algorithm proposed in Section 4.2 to this new dataset, let $\widehat{s}_{i,j}^{\text{new}} \in \{0, 1\}$ denote the estimated state value on the i th day ($i = 1, \dots, \mathcal{N}$). The fractional change in state classification is then defined as

$$\text{Frac. State Classification Changes}(j) \triangleq \frac{\sum_{i=1}^{\mathcal{N}} |\widehat{s}_{i,j}^{\text{new}} - \widehat{s}_i|}{\mathcal{N}}$$

with \widehat{s}_i denoting the state classification on the i th day with the baseline dataset.

In Fig. 10(b), $\text{Frac. Missing Data}(j)$ is plotted as a function of $\text{Frac. State Classification Changes}(j)$. In general, combining terrorism information from two different databases with a clear dichotomy in terms of data-collection standards (criteria for inclusion and non-inclusion of events, source material used, etc.) can introduce a systematic bias in terrorism trends (LaFree and Dugan, 2009). Despite this anomaly, it is clear from Fig. 10(b) that the proposed approach is remarkably robust to a small amount of missing and/or mislabeled data. For example, the addition of an $\approx 5\%$ more data to the baseline dataset leads to essentially no changes in state classification with the baseline dataset. On the other extreme, big additions of even up to 65% more data result in only an $\approx 10\%$ mismatch in state classifications.

7. Concluding Remarks. This work developed a HMM framework to model the activity profile of terrorist groups. Key to this development was the hypothesis that the current activity of the group can be captured completely by certain states/attributes of the group (such as its *Intentions* and its *Capabilities*), instead of the entire past history of the group. In the simplest example of the proposed framework, the group's activity is captured by a $d = 2$ state HMM with the states reflecting a low state of activity (*Inactive*) and a high state of activity (*Active*), respectively. In either state, the days of activity are modeled as a discrete-time Poisson point process with a hurdle-based geometric model being a good-fit for the number of attacks per day. While more general models can be considered, even the simplest framework is sufficient for detecting spurts and downfalls in the activity profile of many groups of interest.

For this purpose, two detection methodologies were proposed in this work. The first of these approaches was a stopping-time based on the EWMA filter that tracks the quantities of interest. This approach does not require knowledge of the underlying model and is hence advantageous. We showed that persistent changes can be detected with this approach. The second approach that requires an estimate of the model parameters is to estimate the most probable state sequence using the Viterbi algorithm. We showed that this approach detects not only non-persistent changes, but also has a low detection delay thus allowing

the identification of key geopolitical undercurrents (or events) that lead to sudden spurts in a group's activity. This early detection of a spurt can then be used for further goals in course of action analysis such as course prediction and policy framing to handle, manage and/or ameliorate the activities of the terrorist group.

Fruitful directions to explore in the future include development of more refined models for the activity profile (such as hierarchical HMMs, switching HMMs, etc.) that incorporate heavy tails and extreme outliers commonly observed in terrorism data. A systematic comparison between the class of self-exciting models and HMMs in explaining the clustering of attacks and a possible bridge between these two classes will also be of interest. In terms of inferencing, non-linear filtering approaches such as particle filters that use *causal* information on terrorist activity and that are computationally feasible will be of importance in practice. Given the intensive nature of data collection and processing that is common for studies of this nature, it would be of interest in developing broad trends and trade-offs in quantitative terrorism studies with a large set of groups from different ideological proclivities.

APPENDIX A: INFORMATION ON MODELS FOR THE NUMBER OF ATTACKS PER DAY STUDIED IN THIS WORK

Table 6 provides information on the density function $P(M_i = k | \mathbf{S}_i = j)$, $k \geq 0$, $j = 0, 1$, the log-likelihood function with observations $x_i \Big|_{i=1}^n$, the ML estimate of the model parameter(s), and the AIC for the six models discussed in Section 3 of the main text. The subsequent sections derive: i) the density function of Y_n and (X_n, Y_n) under the geometric and hurdle-based geometric assumptions on $\{M_i\}$, and ii) Baum-Welch updates for model parameter(s) estimates in these two settings.

A.1. Models for $\{M_i\}$. Table 6 considers six models for $\{M_i\}$ and lists relevant information such as the density function, ML estimate of the parameter with observations $x_i \Big|_{i=1}^n$, AIC, etc. The fits of these models to specific groups such as FARC and Shining Path are studied in the main text of the paper.

It can be seen that the geometric and the Poisson distributions can be obtained as special cases of the Pòlya distribution:

$$\begin{aligned} \text{Pòlya}(r, y) \Big|_{r=1} &= \text{Geometric}(1 - y) \\ \text{Pòlya}\left(r, \frac{\lambda}{\lambda + r}\right) \Big|_{r \rightarrow \infty} &= \text{Poisson}(\lambda). \end{aligned}$$

Further, a straightforward calculation shows that the complementary cumulative distribu-

TABLE 6
Information on observation density models for state j , $j \in \{0, 1\}$, (Support: $k \in \{0, 1, \dots\}$) with
observations $x_i \Big|_{i=1}^n$

Name	Density	Log-likelihood function ($\log(L)$)	ML estimate	AIC
Poisson	$\frac{(\lambda_j)^k e^{-\lambda_j}}{k!}$	$\log(\lambda_j) \cdot \sum_{i=1}^n x_i - n\lambda_j - \sum_{i=1}^n \log(x_i!)$	$\hat{\lambda}_j = \frac{\sum_{i=1}^n x_i}{n}$	$2 - 2\log(L) \Big _{\hat{\lambda}_j}$
Shifted Zipf	$\frac{(1+k)^{-y_j}}{\zeta(y_j)}$	$-n \log(\zeta(y_j)) - y_j \cdot \sum_{i=1}^n \log(1+x_i)$	\hat{y}_j solves $\frac{n\zeta'(y)}{-\zeta(y)} = \sum_{i=1}^n \log(1+x_i)$	$2 - 2\log(L) \Big _{\hat{y}_j}$
Geometric	$(1-\gamma_j) \cdot (\gamma_j)^k$	$n \log(1-\gamma_j) + \log(\gamma_j) \cdot \sum_{i=1}^n x_i$	$\hat{\gamma}_j = \frac{\sum_{i=1}^n x_i}{n + \sum_{i=1}^n x_i}$	$2 - 2\log(L) \Big _{\hat{\gamma}_j}$
Pölya	$\frac{\Gamma(k+r_j)}{\Gamma(k+1)\Gamma(r_j)} \times (1-y_j)^{r_j} (y_j)^k$	$r_j n \log(1-y_j) - n \log(\Gamma(r_j)) + \log(y_j) \cdot \sum_{i=1}^n x_i - \sum_{i=1}^n \log(\Gamma(x_i+1)) + \sum_{i=1}^n \log(\Gamma(x_i+r_j))$	\hat{r}_j and \hat{y}_j solve Eq. 1 and 2 below	$4 - 2\log(L) \Big _{\hat{r}_j, \hat{y}_j}$
Hurdle- based Zipf	$1 - \gamma_j$, if $k = 0$ $\frac{\gamma_j \cdot k^{-y_j}}{\zeta(y_j)}$, if $k \geq 1$	$n_0 \log(1-\gamma_j) - (n-n_0) \log(\zeta(y_j)) + (n-n_0) \log(\gamma_j) - y_j \sum_{i=1}^{n-n_0} \log(x_i)$, where $n_0 = \sum_{i=1}^n \mathbb{1}(\{x_i = 0\})$	$\hat{\gamma}_j = 1 - \frac{n_0}{n}$, \hat{y}_j solves $\frac{\zeta'(y)}{-\zeta(y)} = \frac{\sum_{i=1}^{n-n_0} \log(x_i)}{n-n_0}$	$4 - 2\log(L) \Big _{\hat{\gamma}_j, \hat{y}_j}$
Hurdle- based Geometric	$1 - \gamma_j$, if $k = 0$ $\gamma_j(1-\mu_j) \times (\mu_j)^{k-1}$, if $k \geq 1$	$n_0 \log(1-\gamma_j) + (n-n_0) \log(\gamma_j) + \left(\sum_{i=1}^n x_i - n + n_0 \right) \log(\mu_j) + (n-n_0) \log(1-\mu_j)$	$\hat{\gamma}_j = 1 - \frac{n_0}{n}$, $\hat{\mu}_j = 1 - \frac{n-n_0}{\sum_{i=1}^n x_i}$	$4 - 2\log(L) \Big _{\hat{\gamma}_j, \hat{\mu}_j}$
	Eq. 1: $y = \sum_{i=1}^n x_i / \left(\sum_{i=1}^n x_i + rn \right)$, Eq. 2: $-n \log(1-y) = \sum_{i=1}^n \frac{\Gamma'(x_i+r)}{\Gamma(x_i+r)} - n \frac{\Gamma'(r)}{\Gamma(r)}$, where $\zeta'(\cdot)$ and $\Gamma'(\cdot)$ denote the derivatives of $\zeta(x)$ and $\Gamma(x)$, respectively, and $\mathbb{1}(\cdot)$ denotes the indicator function of the set under consideration.			

tion functions (CCDF) for these six models are respectively given by:

$$\begin{aligned}
P(X > x) \Big|_{\text{Poisson}} &\sim \frac{(\lambda_j)^x}{x!} \cdot (1 - e^{-\lambda_j}) \text{ as } x \rightarrow \infty \\
P(X > x) \Big|_{\text{Shifted Zipf}} &\sim \frac{x^{-(y_j-1)}}{(y_j-1) \cdot \zeta(y_j)} \text{ as } x \rightarrow \infty \\
P(X > x) \Big|_{\text{Geometric}} &\sim (\gamma_j)^x \text{ as } x \rightarrow \infty \\
P(X > x) \Big|_{\text{Pölya}} &\sim \frac{(1-y_j)^{r_j}}{\log\left(\frac{1}{y_j}\right) \cdot r_j - 1!} \cdot (y_j)^x \cdot (x+r_j-1)^{r_j-1} \text{ as } x \rightarrow \infty \\
P(X > x) \Big|_{\text{Hurdle-based Zipf}} &\sim \frac{\gamma_j}{\zeta(y_j) \cdot (y_j-1)} \cdot x^{-(y_j-1)} \text{ as } x \rightarrow \infty \\
P(X > x) \Big|_{\text{Hurdle-based Geometric}} &\sim \frac{\gamma_j}{\mu_j} \cdot (\mu_j)^x \text{ as } x \rightarrow \infty.
\end{aligned}$$

The above list shows that, of the six models studied here, only the shifted Zipf and the hurdle-based Zipf accommodate heavy tails, whereas the other⁸ four models have exponentially decaying tails.

A.2. Density function of Y_n and (X_n, Y_n) for the geometric model. We now derive the probability density function of Y_n given that $\{M_i\}$ is i.i.d. geometric with parameter γ_j in state j ($j = 0, 1$) as in Table 6. Since $Y_n = \sum_{i \in \Delta_n} M_i$, the probability of observing r attacks in the n th time-window with M_i attacks on the i th day is given by

$$(A.1) \quad \prod_{i=1}^{\delta} [(1 - \gamma_j)(\gamma_j)^{M_i}] = (1 - \gamma_j)^{\delta} \cdot (\gamma_j)^{\sum_{i \in \Delta_n} M_i} = (1 - \gamma_j)^{\delta} \cdot (\gamma_j)^r,$$

which is dependent on $\{M_i\}$ only through $\sum_i M_i$. Thus,

$$(A.2) \quad \begin{aligned} P(Y_n = r) &= (1 - \gamma_j)^{\delta} \cdot (\gamma_j)^r \cdot \mathbf{N} \\ &= (1 - \gamma_j)^{\delta} \cdot (\gamma_j)^r \cdot \binom{\delta + r - 1}{r} \end{aligned}$$

where $\mathbf{N} = \binom{\delta + r - 1}{r}$ is the multinomial coefficient that corresponds to the number of ways in which r attacks can be distributed over δ days.

Alternately, $P(Y_n = r)$ for $r \geq 1$ can be derived as

$$(A.3) \quad \begin{aligned} P(Y_n = r) &= \sum_{k=1}^{\min(r, \delta)} P(Y_n = r, X_n = k) \\ &= \sum_{k=1}^{\min(r, \delta)} P(X_n = k) \cdot P(Y_n = r | X_n = k) \\ &= \sum_{k=1}^{\min(r, \delta)} \binom{\delta}{k} (1 - \gamma_j)^{\delta - k} (\gamma_j)^k \cdot P(Y_n = r | X_n = k), \end{aligned}$$

where the upper limit for k in the first line comes from noting that $Y_n \geq X_n$ and $X_n \leq \delta$ and the lower limit follows since $r \geq 1$, and $P(X_n = k)$ in the second line follows a standard binomial density. Note that $P(Y_n = r | X_n = k)$ for $k \geq 1$ is equivalent to computing the probability of arranging $r - k$ balls in k boxes since each attack day has (by definition) at

⁸Note that the Pölya decays exponentially as $x \rightarrow \infty$.

least one attack. This probability is given by

$$\begin{aligned}
 \text{P}(Y_n = r | X_n = k) &\equiv \text{P}\left(\sum_{\ell=1}^k M_{i_\ell} = r \mid \{M_{i_\ell} > 0, \ell = 1, \dots, k\}\right) \\
 &= (1 - \gamma_j)^k (\gamma_j)^{r-k} \cdot \binom{(r-k) + k - 1}{r-k} \\
 &= (1 - \gamma_j)^k (\gamma_j)^{r-k} \cdot \binom{r-1}{r-k},
 \end{aligned}
 \tag{A.4}$$

where the second line follows the same logic as in (A.1) and (A.2). Combining (A.3) and (A.4), we get the following:

$$\begin{aligned}
 \text{P}(Y_n = r) &= \sum_{k=1}^{\min(r, \delta)} \binom{\delta}{k} (1 - \gamma_j)^{\delta-k} (\gamma_j)^k \cdot (1 - \gamma_j)^k (\gamma_j)^{r-k} \cdot \binom{r-1}{r-k} \\
 &= (1 - \gamma_j)^\delta (\gamma_j)^r \cdot \sum_{k=1}^{\min(r, \delta)} \binom{\delta}{k} \cdot \binom{r-1}{r-k}.
 \end{aligned}
 \tag{A.5}$$

It can be seen that (A.5) is consistent with the computation in (A.2) by noting that $\binom{\delta+r-1}{r}$ is the coefficient of x^r in the binomial expansion of $(1+x)^{\delta+r-1}$, whereas $\sum_{k=1}^{\min(r, \delta)} \binom{\delta}{k} \cdot \binom{r-1}{r-k}$ is the coefficient of x^r in the expansion of $(1+x)^\delta \cdot (1+x)^{r-1}$. The constraints $1 \leq k \leq \min(r, \delta)$ follow since $0 \leq r-k \leq r-1$ and $0 \leq k \leq \delta$ in the binomial expansions. Similarly, we can also write the joint density as

$$\text{P}(Y_n = r, X_n = k) = (1 - \gamma_j)^\delta (\gamma_j)^r \cdot \binom{\delta}{k} \cdot \binom{r-1}{r-k}, \quad r \geq k.$$

A.3. Derivation of Baum-Welch estimate for γ_j . The first step in the Baum-Welch algorithm is the verification of the log-concavity of the density function. For this, note that X_n which is Poisson distributed satisfies the log-concavity condition straightforwardly (Juang and Rabiner, 1990, Sec. III) and thus, the Viterbi algorithm can be directly applied to X_n with $\hat{\gamma}_j$. For Y_n , note that

$$\log(\text{P}(Y_n = r)) = \delta \log(1 - \gamma_j) + r \log(\gamma_j) + \log\left(\binom{\delta+r-1}{r}\right),$$

which can be easily checked to be concave in γ_j for all $r \geq 0$. Since the joint density of (X_n, Y_n) has the same form as the density of Y_n in terms of the parameter γ_j , the log-concavity condition also applies to the joint density.

We now develop an expression for the estimate of γ_j (denoted as $\hat{\gamma}_j$) when Baum-Welch algorithm (Rabiner, 1989) is applied to a training-set of N observations $\mathcal{O} = \{O_n, n = 1, \dots, N\}$. Let the corresponding hidden states be $\mathcal{S} = \{S_n, n = 1, \dots, N\}$. The Baum

auxillary function with current/initial estimate of HMM parameters $\hat{\lambda}$ (as a function of the optimization variable λ), $Q(\lambda, \hat{\lambda})$, is given as:

$$Q(\lambda, \hat{\lambda}) = \sum_{\mathcal{S}} \log \left(P(\mathcal{O}, \mathcal{S}|\lambda) \right) \cdot P(\mathcal{O}, \mathcal{S}|\hat{\lambda}).$$

We proceed via the same approach elucidated in (Bilmes, 1998) where the component of the auxillary function due to the observation density is expressed as:

$$\begin{aligned} \frac{Q(\lambda, \hat{\lambda}) \Big|_{\text{Obs. density}}}{P(\mathcal{O}|\hat{\lambda})} &= \frac{\sum_{\mathcal{S}} \log (P(\mathcal{O}, \mathcal{S}|\lambda)) \cdot P(\mathcal{O}, \mathcal{S}|\hat{\lambda}) \Big|_{\text{Obs. density}}}{P(\mathcal{O}|\hat{\lambda})} \\ &= \frac{\sum_{\mathcal{S}} \sum_{n=1}^N \log (P(O_n|S_n, \lambda)) \cdot P(\mathcal{O}, \mathcal{S}|\hat{\lambda})}{P(\mathcal{O}|\hat{\lambda})} \\ &= \frac{\sum_{n=1}^N \log (P(O_n|S_n, \lambda)) \cdot P(\mathcal{O}, S_n|\hat{\lambda})}{P(\mathcal{O}|\hat{\lambda})} \\ (A.6) \quad &= \sum_{n=1}^N \sum_{j=0}^1 \log (P(O_n|S_n = j, \lambda)) \cdot \underbrace{P(S_n = j|\mathcal{O}, \hat{\lambda})}_{\triangleq \gamma_n(j)}, \end{aligned}$$

where the iterative update for $\gamma_n(j)$ follows from (Rabiner, 1989, Sec. IIIA and B). Setting the observations $\mathcal{O} = \{X_n\}$ and using the Poisson density model for the observations in (A.6), a straightforward optimization shows that

$$(A.7) \quad \hat{\gamma}_j = \frac{\sum_{n=1}^N X_n \gamma_n(j)}{\delta \cdot \sum_{n=1}^N \gamma_n(j)}.$$

Similarly, using the density model for $\mathcal{O} = \{Y_n\}$ from Section A.2 in (A.6), we have the following estimate:

$$(A.8) \quad \hat{\gamma}_j = \frac{\sum_{n=1}^N Y_n \gamma_n(j)}{\sum_{n=1}^N (Y_n + \delta) \gamma_n(j)}.$$

Recall that the ML estimate of γ_j from Table 6 with $\{M_i, i \in \Delta_n\}$ as observations is

$$\hat{\gamma}_j \Big|_{\text{ML}} = \frac{\sum_{i \in \Delta_n} M_i}{\sum_{i \in \Delta_n} M_i + \delta} = \frac{Y_n}{Y_n + \delta}.$$

Thus, the Baum-Welch estimate in (A.8) can be understood as a weighted version of the ML estimate with the weights being induced by $\gamma_n(j)$.

Following the same approach as above, it can be seen that with $\mathcal{O} = \{(X_n, Y_n)\}$, the same expression as in (A.8) holds. In other words, with the geometric model for $\{M_i\}$, X_n does not add new information if Y_n is known *a priori* and thus $\hat{\gamma}_j$ does not change. This conclusion should not be surprising as the geometric density is a one-parameter model.

A.4. Extension to hurdle-based geometric model. As noted in equation (4.4) of the main text, X_n is Poisson with parameter $\delta\gamma_j$ in state j where

$$\gamma_j \triangleq \mathbb{P}(M_i > 0 | \mathbf{S}_i = j), \quad j = 0, 1.$$

Following the approach in Section A.2, with $\{X_n, n = 1, \dots, N\}$ as the training-set, γ_j is estimated by the same formula as (A.7) in the geometric case. Note that while $\{X_n\}$ can be used to estimate γ_j , other parameters that describe the density function of M_i cannot be learned with $\{X_n\}$ alone.

We now consider more general models for $\{M_i\}$ that potentially allow $\{Y_n\}$ or the joint sequence $\{(X_n, Y_n)\}$ to result in better inferencing on the states than $\{X_n\}$. In this general setting, the density function of Y_n depends on $\{M_i\}$ not only through $\sum_i M_i$, but also on the precise arrangement of attacks over the δ -day period. Thus estimating the other parameters using either $\{Y_n\}$ or $\{(X_n, Y_n)\}$ becomes difficult, if not impossible. Nevertheless, under the special case of a hurdle-based geometric model (a simple two-parameter extension of the geometric model) elucidated in Table 6, it is possible to estimate γ_j and μ_j using the joint observation sequence $\{(X_n, Y_n)\}$. Following the same steps as in Section A.2, the joint density can be written as

$$\mathbb{P}(Y_n = r, X_n = k) = \binom{\delta}{k} \binom{r-1}{r-k} \cdot (1 - \gamma_j)^{\delta-k} (\gamma_j)^k \cdot (1 - \mu_j)^k (\mu_j)^{r-k}.$$

This density function is log-concave in the parameters and hence the Baum-Welch estimates can be obtained as follows:

$$\begin{aligned} \hat{\gamma}_j &= \frac{\sum_{n=1}^N X_n \gamma_n(j)}{\delta \cdot \sum_{n=1}^N \gamma_n(j)} \\ \hat{\mu}_j &= \frac{\sum_{n=1}^N (Y_n - X_n) \gamma_n(j)}{\sum_{n=1}^N Y_n \gamma_n(j)}, \end{aligned}$$

where the iterative update for $\gamma_n(j)$ is obtained from (Rabiner, 1989, Sec. IIIA and B) using the joint density of $\{(X_n, Y_n)\}$.

APPENDIX B: BACKGROUND INFORMATION ON FARC AND SHINING PATH

This supplementary section motivates the choice of the terrorist groups and the corresponding time-periods of interest that are the focus of this work.

B.1. FARC. The choice of FARC has been motivated by the following reasons:

- Among all terrorist groups active in Colombia over the 1998 to 2006 time-period, FARC is the most dominant with few groups rivalling FARC in terms of strength.
- Being a group of the Marxist-Leninist proclivity, FARC leaves a tell-tale signature in terms of its attack profile. For example, in the 1998 to 2006 time-period that encompassed a total of 604 incidents from the RDWTI database, 512 incidents (or 84.77% of all incidents) could be classified as either an incident on the establishment or infrastructure.

- Over the same time-frame, the competing group Ejército de Liberación Nacional (ELN) could not consolidate its military prowess, was geographically isolated and hence remained sub-dominant relative to FARC. ELN’s categorical rejection of any participation in the drug economy (which is posited to have been responsible for its poor strength) and the smaller demilitarized zone secured by it in negotiations with the government provide further evidence to these claims.
- FARC’s signature of attacks and ELN’s relative dormancy and geographical isolation allows us to classify ambiguous terrorism incidents (on the establishment or infrastructure) as FARC activities with a high degree of confidence.

The 1998 to 2006 time-period is chosen because of the following reasons. In addition to having a clear record of FARC’s activity profile over this period, we are aware of two clear and well-defined *geo-political undercurrents* (events) that have had a major impact on FARC’s activity profile over this period:

- **Event 1:** With Colombia becoming the leading cultivator ([World Drug Report, 2010](#), Fig. 21) of coca by 1997, the Colombian government under President Andrés Pastrana and the U.S. government under President William J. “Bill” Clinton instituted a legislation, popularly known as *Plan Colombia*, under which, among other tasks, U.S. aid for counter-narcotics efforts increased significantly. In the ensuing Presidential elections in 2001-02, Álvaro Uribe won the popular mandate on an anti-FARC platform. The impact of *Plan Colombia* and the impending election of President Álvaro Uribe are the primary reasons behind a spurt in the activity profile of FARC in the late-2001 to 2002 time-period.
- **Event 2:** With FARC adapting itself to the new equilibrium established by the institution of *Plan Colombia* by 2003-04, we observe a return to “normalcy” in terms of the activity profile of FARC. However, with elections to Colombia’s Congress and a re-election bid by President Uribe in the mid- to late-2005 to 2006 period ensured that FARC, which had remained relatively passive in this period, attacked with renewed vigor. This is associated with another spurt in the activity profile of FARC for this period. See ([Cragin and Daly, 2004](#); [Beittel, 2010](#)) (and references therein) for a more detailed description of these events and FARC activities.

B.2. Shining Path. The choice of Shining Path has been motivated by the following reasons:

- The period between 1981 and 1996 corresponds to a full-cycle in terms of Shining Path’s evolution. Shining Path’s evolution toward terrorism began with its inception as an ideological pro-leftist outfit in the late-60s followed by capture of the electoral body of some university campuses across Peru by the mid-70s. By the early-80s, Shining Path had consolidated and re-organized itself as a violent outfit. A cycle of strengthening of counter-insurgency operations, loss of resources by Shining Path, re-organization of available resources, and consequent violence resulted in the capture of its primary ideologue and leader (Abimael Guzmán) in 1992. With this key event, Shining Path entered a phase of organizational catastrophe and decay. While the

organization itself continued to exist in name and through splintered sub-groups, it can be reasonably claimed that the original Shining Path outfit lasted only till the mid-90s.

- As with FARC, Shining Path was the most dominant terrorist group active in Peru in the concerned time-period with the other co-existing sub-dominant group being the Movimiento Revolucionario Túpac Amaru (MRTA). Shining Path also left a tell-tale signature in terms of its targets and attack profile. In fact, 75 of the 163 (or 46.01% of all attacks) over this period from RDWTI were attacks on the establishment with the rest of the attacks being attacks on soft targets. Thus, any ambiguity with regard to incomplete data could be resolved easily.

The data over this period shows the following rich array of geo-political undercurrents resulting in either a spurt or a downfall (at different times) in the activity profile:

- ***Rise and stabilization (1981-1986)***: As Shining Path transformed into a violent terrorist group in the early-80s, an initial tepid response from the Government of Peru resulted in a period where Shining Path quickly seized the initiative and established an equilibrium in terms of activity profile.
- ***Resurgence (1987-1990)***: Attacks on the 1985 Presidential elections resulted in massive counter-insurgency operations by the Government of Peru thereby leading to a slight lull in activity in 1985-86. This event is followed by re-organization of the outfit corresponding to a major resurgence in Shining Path's activity profile in early 1987. However, this activity profile stabilized subsequently.
- ***Heightened state of activity (1991)***: The third phase is marked by an excessively violent 1991. By the end of 1991, Shining Path had control of much of the countryside of the center and south of Peru and had a large presence in the outskirts of Lima.
- ***Death and decay (1992-1996)***: As the organization grew in power, a personality cult grew around its leader, Abimael Guzmán and his philosophy of communism stood elevated to the pantheon of Marxist-Leninist-Maoist thought. However, disrespect for the indigenous culture of Peru, its economic policies that were inimical to indigenous interests, and the excessive violence and vengeance unleashed by Shining Path ensured its massive unpopularity among Peruvians. This institutional apathy in not ensuring the support of the local population led to Guzmán's capture in September 1992. Shortly, thereafter, most of the leadership of Shining Path fell as well. With the last vestiges of rebellion flaming out by 1993, the death of Shining Path as a terrorist outfit is confirmed by a record low number of terrorism incidents in this period. See (Cragin and Daly, 2004) for a more detailed description of these events and Shining Path activities.

REFERENCES

- BADDELEY, A., MOLLER, J. and WAAGEPETERSEN, R. (2000). Non- and semi-parametric estimation of interaction in inhomogeneous point patterns. *Statistica Neerlandica* **54** 329-350.
- BEITTEL, J. S. (2010). Colombia: Issues for the Congress. *Congressional Research Service Report for the U.S. Congress*. Available: [Online]. <http://www.fas.org/sgp/crs/row/RL32250.pdf>.

- BELASCO, A. (2006). The cost of Iraq, Afghanistan, and other global war on terror operations since 9/11. *Congressional Research Service Report for the U.S. Congress*. Available: [Online]. <http://fpc.state.gov/documents/organization/68791.pdf>.
- BILMES, J. A. (1998). A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models Technical Report, International Computer Science Institute, Berkeley CA.
- CLAUSET, A. and GLEDITSCH, K. S. (2012). The developmental dynamics of terrorist organizations. *Submitted to American Journal of Political Science*.
- CLAUSET, A., YOUNG, M. and GLEDITSCH, K. S. (2007). On the frequency of severe terrorist events. *Journal of Conflict Resolution* **51** 58-87.
- COX, D. R. and ISHAM, V. (1980). *Point Processes*, 1st ed. Chapman and Hall.
- CRAGIN, K. and DALY, S. A. (2004). *The Dynamic Terrorist Threat: An Assessment of Group Motivations and Capabilities in a Changing World*. RAND Corporation, Santa Monica, CA.
- CRESSIE, N. A. C. (1991). *Statistics for Spatial Data*, 1st ed. Wiley, NY.
- DIGGLE, P. J. (2003). *Statistical Analysis of Spatial Point Patterns*, 2nd ed. Edward Arnold, London.
- DIXON, P. M. (2002). Ripley's K function. in *Encyclopedia of Environmetrics*, (A. H. El-Shaarawi and W. W. Piegorsch, Eds.), Wiley, Chichester **2** 1796-1803.
- DUGAN, L., LAFREE, G. and PIQUERO, A. (2005). Testing a rational choice model of airline hijackings. *Criminology* **43** 1031-1066.
- DURBIN, J. (1973). *Distribution Theory for Tests Based on the Sample Distribution Function, Volume 9 of CBMS-NSF Regional Conference Series in Applied Mathematics*. SIAM Publishers, Philadelphia.
- ENDERS, W. and SANDLER, T. (1993). The effectiveness of antiterrorism policies: A vector autoregression-intervention analysis. *The American Political Science Review* **87** 829-844.
- ENDERS, W. and SANDLER, T. (2000). Is transnational terrorism becoming more threatening? A time-series investigation. *Journal of Conflict Resolution* **44** 307-332.
- HAUGAARD, L., ISACSON, A. and OLSON, J. (2005). Erasing the lines: Trends in U.S. military programs with Latin America Technical Report, Center for International Policy.
- HAWKES, A. G. (1971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika* **58** 83-90.
- ITERATE, International Terrorism: Attributes of Terrorist Events. Available: [Online]. <http://www.icpsr.umich.edu/icpsrweb/ICPSR/studies/07947>.
- JUANG, B. H. and RABINER, L. R. (1990). The segmental K -means algorithm for estimating parameters of hidden Markov models. *IEEE Transactions on Acoustics, Speech and Signal Processing* **38** 1639-1641.
- KIM, H. (2011). Spatio-temporal point process models for the spread of avian influenza virus (H5N1), Technical Report, University of California, Berkeley. Ph.D. dissertation.
- LAFREE, G. and DUGAN, L. (2007). Introducing the Global Terrorism Database. *Terrorism and Political Violence* **19** 181-204.
- LAFREE, G. and DUGAN, L. (2009). Tracking global terrorism trends, 1970 - 2004. Chapter 3 of *To Protect and To Serve: Policing in an Age of Terrorism*, (D. Weisburd et al., Eds.), Springer Science 43-80.
- LAFREE, G., MORRIS, N. A. and DUGAN, L. (2010). Cross-national patterns of terrorism, comparing trajectories for total, attributed and fatal attacks, 1970-2006. *British Journal of Criminology* **50** 622-649.
- LEWIS, E., MOHLER, G. O., BRANTINGHAM, P. J. and BERTOZZI, A. (2011). Self-exciting point process models of civilian deaths in Iraq. *Security Journal*. Available: [Online]. <http://dx.doi.org/10.1057/sj.2011.21>.
- LINDBERG, M. (2010). Factors contributing to the strength and resilience of terrorist groups. *Grupo de Estudios Estrategicos (GEES) Publication*.
- MIDLARSKY, M. I. (1978). Analyzing diffusion and contagion effects: The urban disorders of the 1960s. *The American Political Science Review* **72** 996-1008.
- MIDLARSKY, M. I., CRENSHAW, M. and YOSHIDA, F. (1980). Why violence spreads: The contagion of international terrorism. *International Studies Quarterly* **24** 262-298.
- MOHLER, G. O., SHORT, M. B., BRANTINGHAM, P. J., SCHOENBERG, F. P. and TITA, G. E. (2011). Self-

- exciting point process modeling of crime. *Journal of the American Statistical Association* **106** 100-108.
- MUELLER, J. and STEWART, M. G. (2011). *Terrorism, Security, and Money: Balancing the Risks, Benefits, and Costs of Homeland Security*. Oxford Univ. Press, NY and London.
- OGATA, Y. (1988). Statistical models for earthquake occurrences and residual analysis for point processes. *Journal of the American Statistical Association* **83** 9-27.
- OGATA, Y. (1998). Space-time point process models for earthquake occurrences. *Annals of the Institute of Statistical Mathematics* **50** 379-402.
- PORTER, M. D. and WHITE, G. (2012). Self-exciting hurdle models for terrorist activity. *Annals of Applied Statistics* **6** 106-124.
- RABINER, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE* **77** 257-286.
- RDWTI, RAND Database of Worldwide Terrorism Incidents. Available: [Online].
<http://www.rand.org/nsrd/projects/terrorism-incidents.html>.
- ROBERTS, S. W. (1959). Control chart tests based on geometric moving averages. *Technometrics* **1** 239-250.
- SANTOS, D. N. (2011). What constitutes terrorist network resiliency? *Small Wars Journal* **7**.
- SESHADRI, V., CSORGO, M. and STEPHENS, M. A. (1969). Tests for the exponential distribution using Kolmogorov-type statistics. *Journal of the Royal Statistical Society, Series B (Methodological)* **31** 499-509.
- SRIVASTAVA, M. S. and WU, Y. H. (1997). Evaluation of optimum weights and average run lengths in EWMA control schemes. *Communications in Statistics - Theory and Methods* **26** 1253-1267.
- TARTAKOVSKY, A. G. and VEERAVALLI, V. V. (2005). General asymptotic Bayesian theory of quickest change detection. *SIAM Theory of Probability and its Applications* **49** 458-497.
- TEERAPABOLARN, K. (2012). A pointwise approximation of generalized binomial by Poisson distribution. *Applied Mathematical Sciences* **6** 1095-1104.
- VEEN, A. and SCHOENBERG, F. P. (2006). Assessing spatial point process models using weighted K -functions. in *Case Studies in Spatial Point Process Modeling*, (A. Baddeley et al., Eds.), *Lecture Notes in Statistics*, Springer, NY **185** 293-306.
- WORLD DRUG REPORT, (2010). Technical Report, United Nations Office on Drugs and Crime.

ADDRESS OF THE FIRST AND THIRD AUTHORS
 DEPARTMENT OF MATHEMATICS
 3620 SOUTH VERMONT AVENUE
 LOS ANGELES, CA 90089
 E-MAIL: vasanthan-raghavan@ieee.org
tartakov@usc.edu

ADDRESS OF THE SECOND AUTHOR
 INFORMATION SCIENCES INSTITUTE
 4676 ADMIRALTY WAY
 MARINA DEL RAY, CA 90292
 E-MAIL: galstyan@isi.edu